

AD A 050887

ARI TECHNICAL REPORT
TR-78-A4

Application of Adaptive Decision Aiding Systems to Computer-Assisted Instruction: Experimental Studies

by

Donald M. May, William H. Crooks, Amos Freedy

PERCEPTRONICS
8271 VARIEL AVENUE
WOODLAND HILLS, CALIFORNIA 91367

FEBRUARY 1978

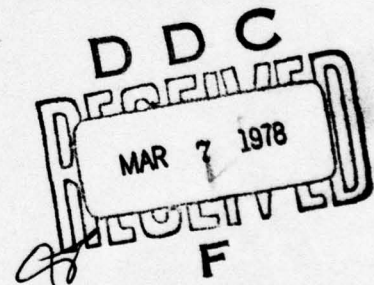
Contract DAHC 19-75-C-0013

Prepared for



U.S. ARMY RESEARCH INSTITUTE for the BEHAVIORAL & SOCIAL SCIENCES
and Defense Advanced Research Projects Agency

Approved for public release; distribution unlimited.



DDC FILE COPY

U. S. ARMY RESEARCH INSTITUTE FOR THE BEHAVIORAL AND SOCIAL SCIENCES

A Field Operating Agency under the Jurisdiction of the
Deputy Chief of Staff for Personnel

J. E. UHLANER
Technical Director

W. C. MAUS
COL, GS
Commander

Research accomplished
under contract to the Department of the Army

PERCEPTRONICS, INC.

NOTICES

DISTRIBUTION: Primary distribution of this report has been made by ARI. Please address correspondence concerning distribution of reports to: U. S. Army Research Institute for the Behavioral and Social Sciences, ATTN: PERI-P, 5001 Eisenhower Avenue, Alexandria, Virginia 22333.

FINAL DISPOSITION: This report may be destroyed when it is no longer needed. Please do not return it to the U. S. Army Research Institute for the Behavioral and Social Sciences.

NOTE: The findings in this report are not to be construed as an official Department of the Army position, unless so designated by other authorized documents.

Unclassified

ARI

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM	
1. REPORT NUMBER Technical Report TR-78-A4	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER rept.	
4. TITLE (and Subtitle) APPLICATION OF ADAPTIVE DECISION AIDING SYSTEMS TO COMPUTER-ASSISTED INSTRUCTION: EXPERIMENTAL STUDIES		5. TYPE OF REPORT & PERIOD COVERED Final 6 Jan 75 - 6 Jan 76	
7. AUTHOR(s) Donald M./May, William H./Crooks Amos/Freedy		14. PERFORMING ORG. REPORT NUMBER PATR-1023-76-1	
9. PERFORMING ORGANIZATION NAME AND ADDRESS Perceptronics, Inc. 6271 Variel Avenue Woodland Hills, CA 91364		15. CONTRACT OR GRANT NUMBER(s) DAHC 19-75-C-0013	
11. CONTROLLING OFFICE NAME AND ADDRESS US Army Research Institute for the Behavioral and Social Sciences (PERI-OK), 5001 Eisenhower Ave., Alexandria, VA 22333		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 20762717A764	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		11. REPORT DATE Feb 78	
		12. NUMBER OF PAGES 66	
		13. SECURITY CLASS (of this report) Unclassified	
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE	
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)			
18. SUPPLEMENTARY NOTES This contract was supported in part by the Department of Defense Advanced Research Projects Agency, 1400 Wilson Boulevard, Arlington, Virginia 22209.			
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Computer-Assisted Instruction Decision Aiding Electronics Training Decision Training Adaptive Training Artificial Intelligence Troubleshooting			
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This report describes a system for Computerized Decision Training (CDT) which focuses on improving and sharpening higher order cognitive skills in judgmental decision making. The CDT system incorporates an adaptive computer program which learns the student's value structure, and uses this structure to train the student in practical decision making. This report describes the application of decision models in training, and presents the features of the CDT system as it is			

DD FORM 1473


EDITION OF 1 NOV 65 IS OBSOLETE

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

20. applied in electronic troubleshooting. Experimental evaluations have demonstrated that the adaptive decision model accurately models the student's performance and that the adaptively-selected instructions sometimes improve troubleshooting performance. 

Unclassified

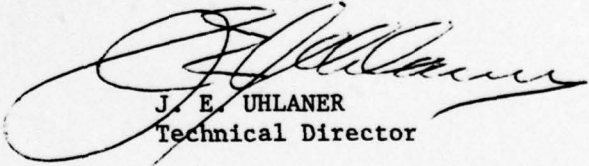
SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

FOREWORD

The Educational Concepts and Evaluation Work Unit Area of the Army Research Institute for the Behavioral and Social Sciences (ARI) performs research and development in areas of educational technology with applicability to military training. Of special interest is research in the area of computer-based training systems. Development and implementation of such systems is seen as a solution to such current Army problems as a shortage of qualified instructor personnel, a student population of widely varying abilities, and increased training costs. Computer-based training systems also provide the potential to increase training effectiveness and efficiency by increasing the extent to which the training process can be made to adapt to the characteristics and performance of the individual student.

This Technical Report describes the second phase of a research effort to develop a technique for individualizing training through the use of "artificial intelligence" techniques. The results of the first phase are documented in an earlier report (May, Crooks, Purcell, Lucaccini, Freedy, and Weltman, 1974). In order to accomplish this research, ARI's resources were augmented by contract with Perceptronics, Inc., an organization selected as having unique capabilities for research and development in this area.

The entire research work unit area is responsive to the requirements of RDT&E Project 2Q762717A764, "Educational and Training Technology," of the 1975 ARI Work Program, and to special requirements of the Product Manager, Computerized Training Systems.


J. E. UHLANER
Technical Director

ACCESSION for	
NTIS	Write Section <input checked="" type="checkbox"/>
DDC	Buff Section <input type="checkbox"/>
UNANNOUNCED	<input type="checkbox"/>
JUSTIFICATION	
by	
DISTRIBUTION/AVAILABILITY CODES	
SPECIAL	
A	

APPLICATION OF ADAPTIVE DECISION AIDING SYSTEMS TO COMPUTER-ASSISTED
INSTRUCTION: EXPERIMENTAL STUDIES

BRIEF

Requirement:

To continue the development and evaluation of a computer-based system which uses adaptive techniques to train electronic troubleshooting procedures. During the training process, the student troubleshoots a simulated electronic circuit by making test measurements and replacing the malfunctioning part. The key component of the system is an adaptive program, based on an Expected Utility (EU) decision model, which "learns" the student's utilities for troubleshooting decisions. These utilities can then be compared with those of an expert, and instructional feedback can be provided to reduce the discrepancies between the two sets of utilities. Specific objectives were: (a) to refine the student/computer interface and computer algorithms; (b) to develop measures of student performance and incorporate them into a student performance summary report; (c) to determine the adequacy of the EU model used to describe student and expert troubleshooting behavior; (d) to evaluate the effects of training with the system on student performance; and (e) to develop a method for providing adaptive instructional feedback to the student.

Procedure:

The student/computer interface was modified to simplify the interaction for the students. Measures of student performance, some of which are "traditional" measures, and others of which are unique to this system, were included as part of a routine student diagnostic report. The utility estimation algorithms were also simplified.

The adequacy of the EU decision model, which is used to describe both student and expert behavior, was evaluated in several ways, first using an electronics expert who diagnosed the circuit using a fixed strategy, then using a "simulated" student, and finally using a group of electronics students.

An experiment was conducted using as subjects eight electronics students who were given 4 1/2 hours training on the system, divided into three sessions. During the second session they were provided with the expert model's probabilities of action outcomes. These probabilities were not available during the first and third sessions.

A method for providing adaptive instructional feedback to the students was developed. The method uses the utilities for "key"

measurements, that is, those measurements identified by an expert as being of critical importance in the troubleshooting process. The specific feedback provided to the student is based on the relative values of the key utilities. A preliminary evaluation of the training effectiveness of the feedback was conducted.

Findings:

Records of all student responses can be recorded for later analysis. A diagnostic report summarizing student performance is printed at the completion of each problem.

The adaptive EU model rank orders the utilities of an expert technician accurately. After practice, it predicts all choices of the simulated student correctly. Approximately 75% of the choices of "real" students are predicted correctly.

The presentation of probabilities of action outcomes improves both student performance and the predictive success of the EU model.

Student performance on the system improves with practice even in the absence of any utility-based feedback.

The presence of utility-based feedback does not consistently modify student performance in the desired fashion.

Utilization of Findings:

These findings provide sufficient evidence to justify continued development and evaluation of the system. The goal of this process is the cost and training effectiveness evaluation of a prototype version within an ongoing course of instruction at an Army school. These findings will be used by ARI, the U.S. Army Training and Doctrine Command Training Support Center, and the U.S. Army Signal School to determine resource requirements for the evaluation and possible future implementation.

TABLE OF CONTENTS

	<u>PAGE</u>
1. SUMMARY	1-1
1.1 Accomplishments	1-3
1.2 System Optimization	1-3
1.3 Research Approach	1-3
1.3.1 First Year Program	1-3
1.3.2 Second Year Program	1-4
2. ADAPTIVE TECHNIQUES IN COMPUTER ASSISTED DECISION TRAINING	2-1
2.1 Background	2-1
2.2 Adaptive Computer Aiding	2-1
2.2.1 Computer Aiding	2-1
2.2.2 Aiding by Means of Decision Analysis	2-2
2.3 Methods of Quantitative Utility Assessment	2-4
3. THE CDT SYSTEM	3-1
3.1 Overview	3-1
3.2 Training Procedure	3-2
3.2.1 Troubleshooting Operations	3-3
3.3 Task Simulator	3-8
3.4 Expert Model	3-10
3.4.1 Information Gain	3-12
3.4.2 Simulated Student	3-14
3.4.3 Requests for Help	3-15
3.5 Student Model and Dynamic Utility Assessment	3-16
3.6 Instructional Logic	3-19
3.7 Diagnostic Report of Student Performance	3-22
3.8 System Software	3-24
3.9 Instructor/Computer Interactions	3-26
4. SYSTEM PERFORMANCE EVALUATION	4-1
4.1 Engineering Evaluation and Analysis	4-1
4.1.1 Training Algorithm Validation	4-1
4.1.2 Estimated Values and Expert Technician Ranking	4-3
4.1.3 Comparison of Student Performance with Model Performance	4-6
4.2 Student Performance Evaluation	4-8
4.3 Evaluation of Instructions	4-14
4.4 Discussion	4-22
5. REFERENCES	5-1

LIST OF FIGURES

<u>Figures</u>	<u>Page</u>
3-1 Student's Display of a Troubleshooting Problem	3-4
3-2 Student's Display of a Partially-Completed Problem	3-6
3-3 Schematic Representation of Dynamic Value Estimator	3-18
4-1 Estimated Values for Target Measurement Outcomes	4-2
4-2 Predictive Success of Decision Model during Simulated Student Performance	4-4
4-3 Relative Utility for Key Measurements as a Function of Fault Problems (Expert Technician)	4-5
4-4 Student and Student-Model Cost Expenditures for Problem Solution	4-7
4-5 Estimated Utilities as a Function of Troubleshooting Decisions (Student A)	4-9
4-6 Estimated Utilities as a Function of Troubleshooting Decisions (Student C)	4-10
4-7 Student Decision Time as a Function of Training	4-12
4-8 Student Decision Efficiency as a Function of Training	4-13
4-9 Predictive Success of Decision Model During Student Training	4-15
4-10 Relative Utility for Key Measurements as a Function of Fault Problems (Student JK)	4-19
4-11 Relative Utility for Key Measurements as a Function of Fault Problems (Student BW)	4-21
4-12 Student Cost Efficiency as a Function of Training	4-22

LIST OF TABLES

<u>Tables</u>		<u>Page</u>
3-1	Performance Summary Indices	3-25
4-1	Utility Training Instructions and Selection Criteria	4-16

GLOSSARY OF SYMBOLS

$\alpha_{i,j}$	Information gain function associated with $X_{i,j}$
B	Lower stopping boundary (Bad)
ϵ	Set membership
EU	Expected Utility
EV	Expected Value
f_j	Faults associated with $X_{i,j}$ divided by the number of faults still possible.
G	Upper stopping boundary (Good)
γ	Value training constant
i	An index over outcomes.
j	An index over measurements
k	An index over faults
λ	Likelihood ratio
$L_{i,j}$	Set of faults corresponding to $X_{i,j}$.
L_x	Set of faults still possible given \bar{X} .
$P_{i,j}$	Same as $P(X_{i,j})$
$P(k), P_k$	Probability of fault K.
$P(M_\ell)$	Failure probability of module ℓ .
$P(X_{i,j})$	Probability of obtaining $X_{i,j}$
$V_{i,j}$	Value or utility for outcome $X_{i,j}$
\bar{X}	Measurement history.
$X_{i,j}$	Outcome i of measurement j.
X_j	Measurement j.
\cap	Intersection

1. SUMMARY

This report describes the results of the second year's effort in the development of a new system for Computerized Decision Training (CDT).¹ The CDT system combines the principles of artificial intelligence, decision theory, and adaptive computer-assisted instruction. Training focuses on higher order cognitive skills in judgmental decision-making. Realistic simulation of training problems permits student application of decision-making skills in real-life contexts, increasing the potential for transfer of training to field situations.

The CDT concept incorporates an adaptive computer program which learns the student's diagnostic and decision value structure, compares this structure to that of an expert, and adapts the instructional sequence to eliminate discrepancies. An expected utility (EU) model of decision making is the basis of the student and instructor models which, in conjunction with the decision task simulator, form the core of the CDT system. The student model is dynamically adjusted using a trainable network technique of pattern classification. Heuristic algorithms generate the training instructions and adjust the problem presentation sequence. The instructor model also generates suggested actions in response to student requests for assistance.

The present training system focuses on electronic troubleshooting. The student's task is to troubleshoot a complex circuit by making various test measurements, replacing the malfunctioning part, and making final verification measurements. The student utilities of interest are those for information gained through the measurements, and for replacement of circuit modules. Troubleshooting provides an excellent application for the CDT methodology because it is heavily dependent on judgment and probabilistic inference. In addition, troubleshooting as such is of great practical importance in numerous Army systems, and it lends itself to economical implementation for training purposes.

¹Special recognition is given to Denis Purcell and Michael Kuppin for their efforts in developing the instruction-selection and experimenter/computer interaction programs.

Work to date has produced an operational system which demonstrates the feasibility of applying artificial intelligence techniques to computer assisted instruction in a minicomputer environment. Experimental evaluations of the CDT system have demonstrated that the adaptive decision model accurately learns the utilities of an expert technician and that students can effectively use the simulated troubleshooting task. Additional software evaluations and developments have optimized the decision model training algorithm, provided a smooth student/computer sequence of interactions, and created a report routine for diagnosing and summarizing student performance.

Research findings include the following:

- (1) The adaptive EU model with information gain accurately learns and rank orders the utilities of an expert technician or student as they interact with the simulated decision tasks.
- (2) Presentation of probabilities of decision alternatives represents a valuable form of decision aid which increases decision consistency, increases troubleshooting speed, and reduces repair cost. Students with little or no prior experience with probability are quickly able to learn to use the action probabilities.
- (3) Performance improves with experience on the system.
- (4) The EU-based model of the student performs better than the student himself when the student's utilities are used in a consistent simulated decision sequence.

1.1 Accomplishments

Research and development activities during 1975 centered on the following objectives: (1) system optimization; (2) engineering analysis and evaluation; (3) development of diagnostic reporting routines for student and model assessment; (4) experimental evaluation of student performance; and (5) development of instructional feedback algorithms.

The following paragraphs provide a summary of the current work tasks and accomplishments.

1.2 System Optimization

The first year of program activities produced a prototype decision training system that includes (1) a circuit simulation with capabilities for CRT display of circuit measurement results, (2) an adaptive decision model of the student, (3) a decision model of an expert technician, and (4) a student/computer interaction sequence that incorporates various instructional activities. System optimization during the current year took two forms: (1) changes in the CDT system which improved the external system characteristics (i.e., the student/computer interaction), and (2) improvements in the system internal functions (i.e., adaptive algorithm parameters).

1.3 Research Approach

1.3.1 First Year Program. The first year's work involved selection of decision task, simulation of the essential decision features of the task, implementation of the adaptive decision model, and development of a student/computer interaction sequence. The result of the year's efforts was an operational prototype simulation trainer.

1.3.2 Second Year Program. The second year's efforts were directed toward (1) optimization of the computer algorithms and student/computer interactions, (2) operational analysis of the adaptive decision model in comparison to consistent decision making by a simulated student and by an expert technician, (3) experimental evaluations of student's performance with the simulated electronic circuit and the associated decision aids, and (4) development of adaptive instructional feedback based on the estimated values in the decision model of the student.

The results of the analyses during the second year have demonstrated that the adaptive decision model tracks the decision performance of a consistent expert technician, and the estimated values accurately reflect the relative ranking of the critical decision alternatives identified by the technician. The analyses have also shown that student technicians use the simulated troubleshooting task and can improve their decision-making efficiency after extended practice with the CDT system.

2. ADAPTIVE TECHNIQUES IN COMPUTER ASSISTED DECISION TRAINING

2.1 Background

A central theme in the field of educational technology is the creation of methods which allow the individualization of instruction. Much work has been done toward this end in the field of Computer Assisted Instruction (CAI). This CAI background relevant to the development of the CDT system is reviewed in our first annual report (May, Crooks, Purcell, Lucaccini, Freedy, and Weltman; 1974). The present report focuses on the background work in computer decision aiding and utility* assessment upon which the CDT system also is based.

2.2 Adaptive Computer Aiding

Computer aiding systems have recently evolved from inflexible routines to adaptive programs capable of high-level interaction and initiative. Adaptive (or intelligent) components are taking over many of what were once considered uniquely human functions, such as learning, pattern recognition, problem solving, and inferential decision making. Employing such functions, an intelligent aiding system can analyze alternative actions, recommend responses, and even perform various tasks autonomously.

2.2.1 Computer Aiding. Computer aiding systems are devices that simplify or otherwise facilitate the performance of some specific task. Applications cover such diverse situations as continuous dynamic control of remotely piloted vehicles, intelligence gathering, information flow in command and control operations and CAI. The emphasis in these advanced

*In keeping with the literature on decision making, the term "utility" is used to denote "subjective value".

systems is on the aiding of decision making processes, including such contributions as data organization and display, establishment of procedures to select courses of action, mathematical optimization (linear programming, optimal control and the like), and decision analysis (Brown, Hoblitzell, Peterson and Ulvila, 1974; Howard, 1968; Nickerson and Fehrer, 1974; Weisbrod, Davis and Freedy, 1975).

2.2.2 Aiding by Means of Decision Analysis. Applying decision analysis to decision aiding involves seven basic steps. These steps cover the relevant aspects of defining decision choices and parameters required to establish a measurable criterion of optimal choice (Payne, Miller, Ronney, 1974):

- (1) Identification of pertinent information.
- (2) Definition of alternatives.
- (3) Definition of structure for related data parameters, events and alternatives.
- (4) Characterization of uncertainty of continuous parameters.
- (5) Estimation of event probabilities.
- (6) Transformation of multi-attribute measures into a single utility for each possible outcome.
- (7) Selection of the best alternative through normative evaluation criteria.

The basic components of the criteria are probabilities and utilities. It is necessary to determine the probabilities of alternative decision outcomes and assess the utilities that the decision maker has for these outcomes. Probabilities can normally be estimated by objective measurement or from prior probabilities elicited from experts. Then the prior probabilities can be aggregated, using Bayesian or probabilistic information processing (Edwards, 1962; Kelly and Peterson, 1971), to obtain posterior probability estimates. These techniques also provide a mechanism for updating the probabilities as new data becomes available.

For the most part, experimental studies of computer-based decision aids have shown that operators are able to function effectively with machine support. Hanes and Gebhard (1966), in a realistic simulation of anti-aircraft warfare, found that Naval commanders freely accepted computer advice in a tactical command action. Similarly, Miller, Kaplan and Edwards (1967) demonstrated the efficiency of combining human value judgment and machine policy selection to perform aircraft dispatching in a tactical air command system. The interaction was found to be superior in performance to unaided human dispatching. A good demonstration of a computer-based information system is found in the U.S. Army's Simulated Tactical Operations System (SIMTOS). SIMTOS is an interactive information system for command and control operations designed to complement the man's information processing and decision making capabilities. The system has been used to test the effects of a number of procedural and information control factors on system performance (Baker, 1970, 1974).

On the other hand, some of the early work in computer aiding engendered a competitive situation between man and machine, since the machine acted as a "surrogate" or replacement of the operator (Vaughn and Mavor, 1972). The trend toward adaptive, interactive systems has ameliorated some of the problems of conflict by emphasizing the man-with-a-computer concept rather than the man-versus-computer form. The more interactive "staff" functions are characterized as providing requested information, suggesting alternatives, alerting the operator to important data, and performing other decision subtasks in an advisory manner (Vaughn and Mavor, 1972; Halpin, Thornberry and Streufert, 1973).

Previous work (Davis, Weisbrod, Freedy, and Weltman, 1975) using an adaptive aiding system called ADDAM concluded the following:

- (1) The ADDAM program adaptively estimates operator utilities in realistic decision-making situations.

- (2) Utility estimation is consistent over subsets of the total outcome set.
- (3) Utility estimation rapidly stabilizes for consistent operator decision behavior.
- (4) Decision recommendations based on adaptive utility estimates are well accepted by experienced operators.
- (5) Availability of individualized recommendations markedly improves decision-making performance by (a) allowing the individual operator to maintain near-maximum expected utility; and (b) reducing variability among different operators.

Of particular importance for practical use of adaptive aiding is conclusion (2), which indicates that a large set of utilities can be trained in trials involving only a small number of utilities at a time, and conclusion (3) which indicates that utilities may be estimated in a time period which is quite reasonable for many simulated and operational decision-making tasks.

The CDT system concept grew out of computer aiding. Some aspects of aiding are used directly in the CDT system such as presentation of action probabilities, presenting measurement results in a semi-interpreted form, and the HELP function which provides suggested next actions and a list of remaining alternatives. The ability of the CDT system to track the student, diagnose deficiencies, and give instructional feedback is closely related to similar aiding functions.

2.3 Methods of Quantitative Utility Assessment

Reliable methods for quantitative assessment of utilities are a major area of difficulty. A number of techniques have been suggested and used, both in the research literature (Kneppreth, Gustafson, Leifer, and Johnson, 1974),

and in the emerging discipline of applied decision analysis (Brown, Hoblitzell, Peterson and Ulvila, 1974). The problem of assessing utilities has become especially acute in recent years because of the growing interest in quantitative decision analysis.

A large number of techniques for utility assessment have been suggested. These may be classified according to the measurement and computational processes used to estimate them. Comprehensive reviews of utility assessment techniques have been prepared (e.g., Kneppreth, Gustafson, Leifer, and Johnson, 1974; Fishburn, 1967). Three major classes of utility assessment techniques are briefly reviewed here:

- (1) Elicitation of utilities through direct judgment. An analyst asks the DM directly to give his value for each decision outcome. These values are normally measured as point values for a particular outcome. These values can be obtained directly using a wholistic approach (Beach, 1973). However, since outcomes usually have several attributes they are often decomposed into single attribute outcomes. The single attribute utilities thus elicited are then combined linearly, to yield the DM's utilities for the more complex outcomes.
- (2) Inference from behavior in simple gambles and other decision games. This technique requires the DM to respond to a series of simple gambles or decision games which usually involve financial reward. The DM's choices form a data base from which his utilities are inferred, usually by indifference techniques. These techniques have been used by a number of investigators (e.g., Tversky, 1967), but are mainly constrained to laboratory/research settings.
- (3) Dynamic estimation through decision observation. This approach calls for direct observation of decision behavior in real-world or simulated real-world contexts. A primary example of this approach is the ADDAM System (Freedy, Weisbrod, Davis, May, and Weltman, 1974) and the CDT system described herein. Other attempts

to use this approach are more concerned with modeling the decision maker's gross behavior while determining his utilities. One good example is the bootstrapping technique of Dawes (1970) which uses a brute force linear model.

The advantages of the Dynamic Observation technique are as follows:

(1) Utilities are estimated non-verbally, without the need for a skilled analyst highly trained in utility estimation techniques. Indeed, the DM need not be aware that his utilities are being assessed. Utilities can be estimated rapidly and the technique is not limited by the number of possible decision outcomes. (2) The utilities are measured on a common scale and are combinatory. (3) The utility assessment technique responds adaptively to changes in values and the utilities are automatically validated by direct comparison with the DM's real-world behavior. These advantages have important implications for adaptive decision training.

The adaptive technique assumes an expected utility maximization paradigm for modeling decision behavior, and uses a pattern recognition algorithm to successively adjust the model to fit observed decision behavior. The underlying expected utility (EU) model assumes that the operator chooses that action whose expected (probability weighted) utility of outcome is highest (Krantz, Luce, Suppes and Tversky, 1971). EU models, of course, are not a panacea for structuring decision models. Lichtenstein and Slovic (1971) argue that descriptive models must take cognitive factors into account, Luce and Suppes (1965) question the use of deterministically maximized choices rather than stochastic choices, and Wendt (1970) and Coombs and Pruitt (1960) contend that the EU model should be modified to account for preferences in variance of outcome. In general, though, the usefulness of EU model is conceded in situations where the number of choices is low and the decision maker can relate to all attributes in terms of probabilities (Goodman, Saltzman, Edwards, and Krantz, 1971). Also, the EU models have the advantage of modeling both descriptive and normative (optimal) behavior, unlike most of the heuristic-based models (Wendt, 1973).

Because the utility estimator is being continuously adjusted it is useful to examine the behavior of the utilities under various conditions. If the probability patterns are linearly separable into categories (decisions), the utility estimator will learn to classify them perfectly after a finite number of steps. Since adjustment takes place only when there are classification errors, each utility will converge to a single value. If the operator's values change, the utility estimator will begin making errors again, adjustment will take place, and the utilities will converge to a new set of values.

If the patterns are not linearly separable, a different situation arises since the utility estimation can never learn to classify perfectly. In a conventional pattern classifier, linear inseparability is reflected in the error rate. In a system which is continuously being adjusted, this error rate keeps the utilities from converging to a single value. However, the utilities may approach a steady state value within a range of variance.

The accuracy in prediction of behavior and the degree to which the utility estimates converge can be used as measures of the validity of the utility estimates. Perfect predictive validity would result in the convergence of the utilities to a single set of values. However, given the limitations of human memory and information processing, and the inconsistencies in human behavior, it would be unreasonable to expect the EU model to be perfectly predictive. If the operator's behavior is consistent with the model "most of the time", the steady state variability of the utilities will be small. Likewise, if his behavior is "erratic" it will be large.

The validity of the utility estimating approach has been substantiated both in preliminary experiments in this program and by experiments in two completely separate simulations: (1) An intelligence gathering task in a simulated dynamic environment (Weisbrod, Davis, Freedy, and Weltman, 1974); and (2) An automated control allocation experiment (Steeb, Artot, Crooks, Freedy, and Weltman, 1976). In the first simulation the validity of the utility estimation technique was tested by comparison of predicted operator decision

behavior with actual behavior and by comparison with directly expressed preferences. A small sample of three subjects showed that a derived utility maximizing model predicted more than 95% of the deployments of intelligence sensors actually made by the subjects performing an intelligence-gathering task. Similarly, in a subsequent sample of nine subjects, a high correlation ($r=.82$, $p<.01$) was found between estimated and operator-expressed sensor preferences.

3. THE CDT SYSTEM

3.1 Overview

In essence, the CDT system is a computer program with the ability to simulate or model three main elements in the diagnostic decision training situation. These elements are:

- (1) Decision Making Task. The program simulates the particular task environment with which the student interacts in the performance of the task. Numerous examples of real life problems can be simulated using the CDT approach.
- (2) An Expert. Part of the program acts like an "expert" who serves as both a source of help to the student, and as a standard against which the student is judged.
- (3) The Student. A separate part of the program constantly monitors the student's decision making behavior, and constructs a model of the student similar in form to the simulated expert.

The CDT system uses an Expected Utility (EU) model of the student and of the instructor. In the student model, the EU principle is used as a basis for dynamically estimating the student's utilities for action outcomes. In the instructor model, the EU principle is used to generate suggested actions in response to student requests for assistance. The student's utilities for outcomes, in comparison with the expert's utilities for the outcomes, provide the framework for generating remedial instructions. The utility of an outcome provides a measure of its worth or relative real world desirability.

The EU model is used to define instructor and student choice behavior in selecting courses of action in a diagnosis and decision task. The

probabilities of outcomes make up a measure of diagnosis while the utilities provide a measure of the relative real world desirability or worth of an outcome. In the troubleshooting context, probabilities are associated with the likelihood of occurrence of measurement outcomes and circuit module faults as inferred from observed symptoms. The symptoms define the information available about the equipment at a given time. The utilities are associated with the worth of knowledge about certain action outcomes, and the contribution of this knowledge to determining technical circuit problems. In essence, the expected utility model defines the relative desirability of performing a certain measurement or part replacement under a given set of circuit symptoms. The model is expressed in terms of probabilities of obtaining an outcome -- such as a certain measurement result -- and the relative value for the information about this measurement. The instructor utilities are calculated using the adaptive utility estimation technique and then stored in a utility matrix. These instructor utilities are thus available during the training task to serve as a standard against which the estimated student values are compared. The aggregated probabilities of the instructor model are displayed to the student. Rather than using a static set of utilities, as in the instructor model, the student utilities are dynamically adjusted throughout training using the adaptive, on-line utility adjustment subprogram.

3.2 Training Procedure

Training in the CDT system is provided by three procedures designed to introduce the student to the system, provide practice in decision making, supply instructions to direct the decision making, and assist the student when difficulties are encountered.

Briefing. The student's first encounter with the decision training system involves a description of the system and training objectives and a screening process. During the introductory briefing the student is told what will be expected of him and how his performance will be assessed. Description

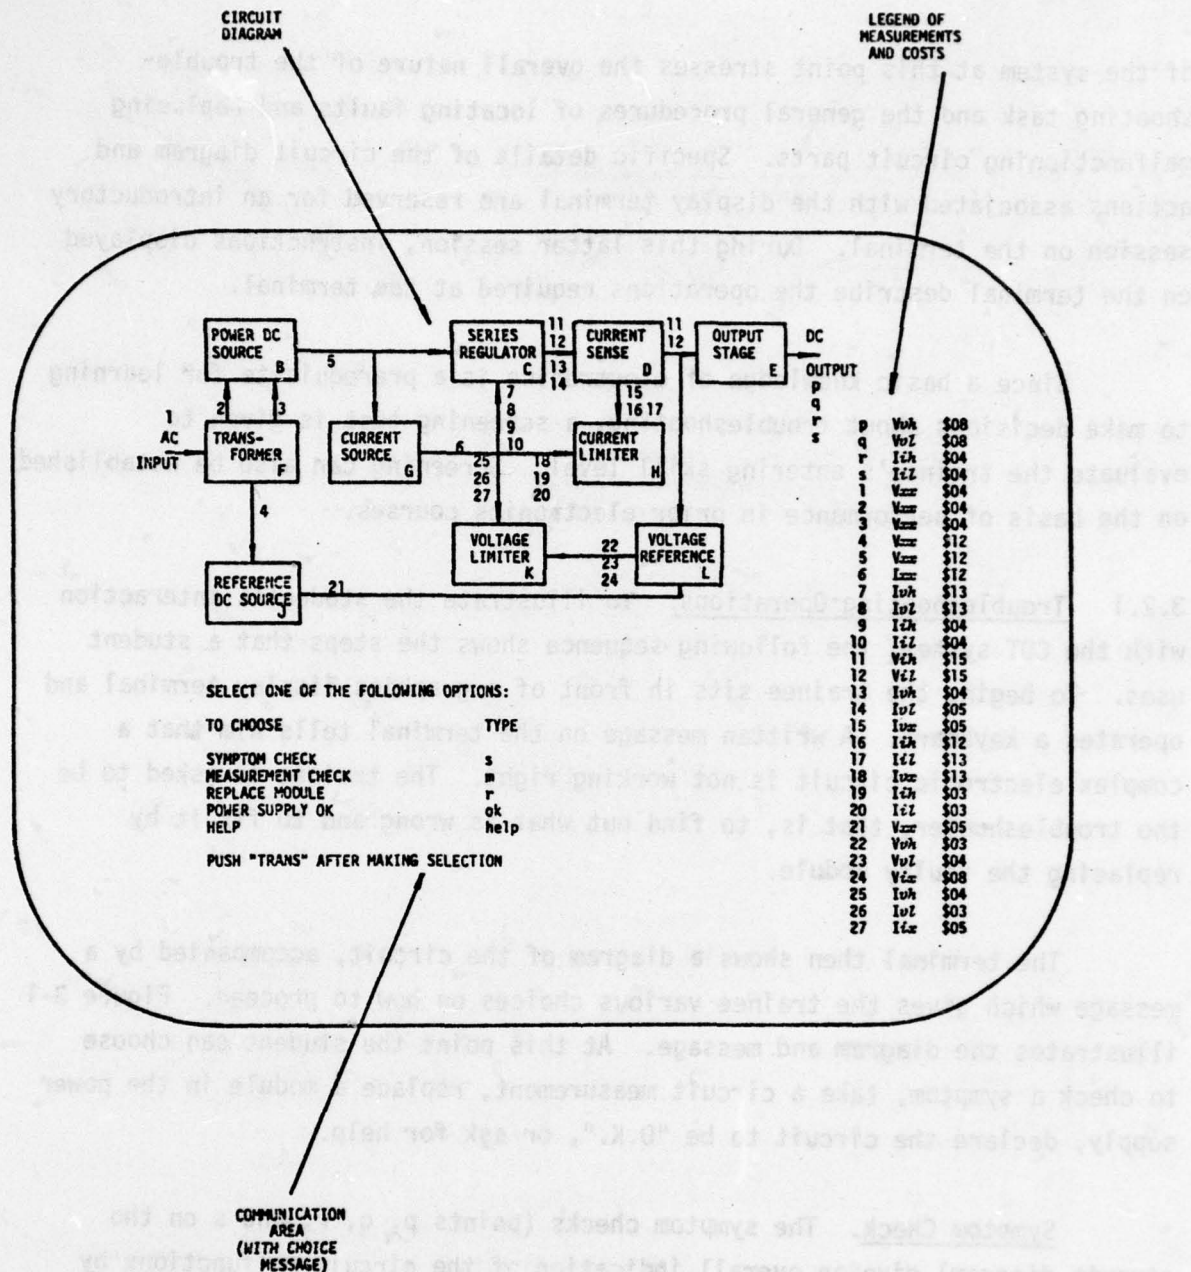
of the system at this point stresses the overall nature of the troubleshooting task and the general procedures of locating faults and replacing malfunctioning circuit parts. Specific details of the circuit diagram and actions associated with the display terminal are reserved for an introductory session on the terminal. During this latter session, instructions displayed on the terminal describe the operations required at the terminal.

Since a basic knowledge of electronics is a prerequisite for learning to make decisions about troubleshooting, a screening test is given to evaluate the trainee's entering skill level. Screening can also be established on the basis of performance in prior electronics courses.

3.2.1 Troubleshooting Operations. To illustrate the student's interaction with the CDT system, the following sequence shows the steps that a student uses. To begin, the trainee sits in front of a graphics display terminal and operates a keyboard. A written message on the terminal tells him that a complex electronic circuit is not working right. The trainee is asked to be the troubleshooter, that is, to find out what is wrong and to fix it by replacing the faulty module.

The terminal then shows a diagram of the circuit, accompanied by a message which gives the trainee various choices on how to proceed. Figure 3-1 illustrates the diagram and message. At this point the student can choose to check a symptom, take a circuit measurement, replace a module in the power supply, declare the circuit to be "O.K.", or ask for help.

Symptom Check. The symptom checks (points p, q, r, and s on the circuit diagram) give an overall indication of the circuit malfunctions by measuring the current or voltage output of the power supply. The output voltage can be measured with the voltage control on the power supply set to the high state (symptom p) or to the low state (symptom q). Similarly, the



**STUDENT'S DISPLAY
OF A TROUBLESHOOTING PROBLEM**

FIGURE 3-1

output current can be measured with the current control set to the high state (symptom r) or to the low state (symptom s).

Measurement Check. The student can take a circuit measurement to gain information about the internal behavior of the circuit. Thus, he can isolate malfunctions to specific modules by measuring the inputs and outputs of those modules.

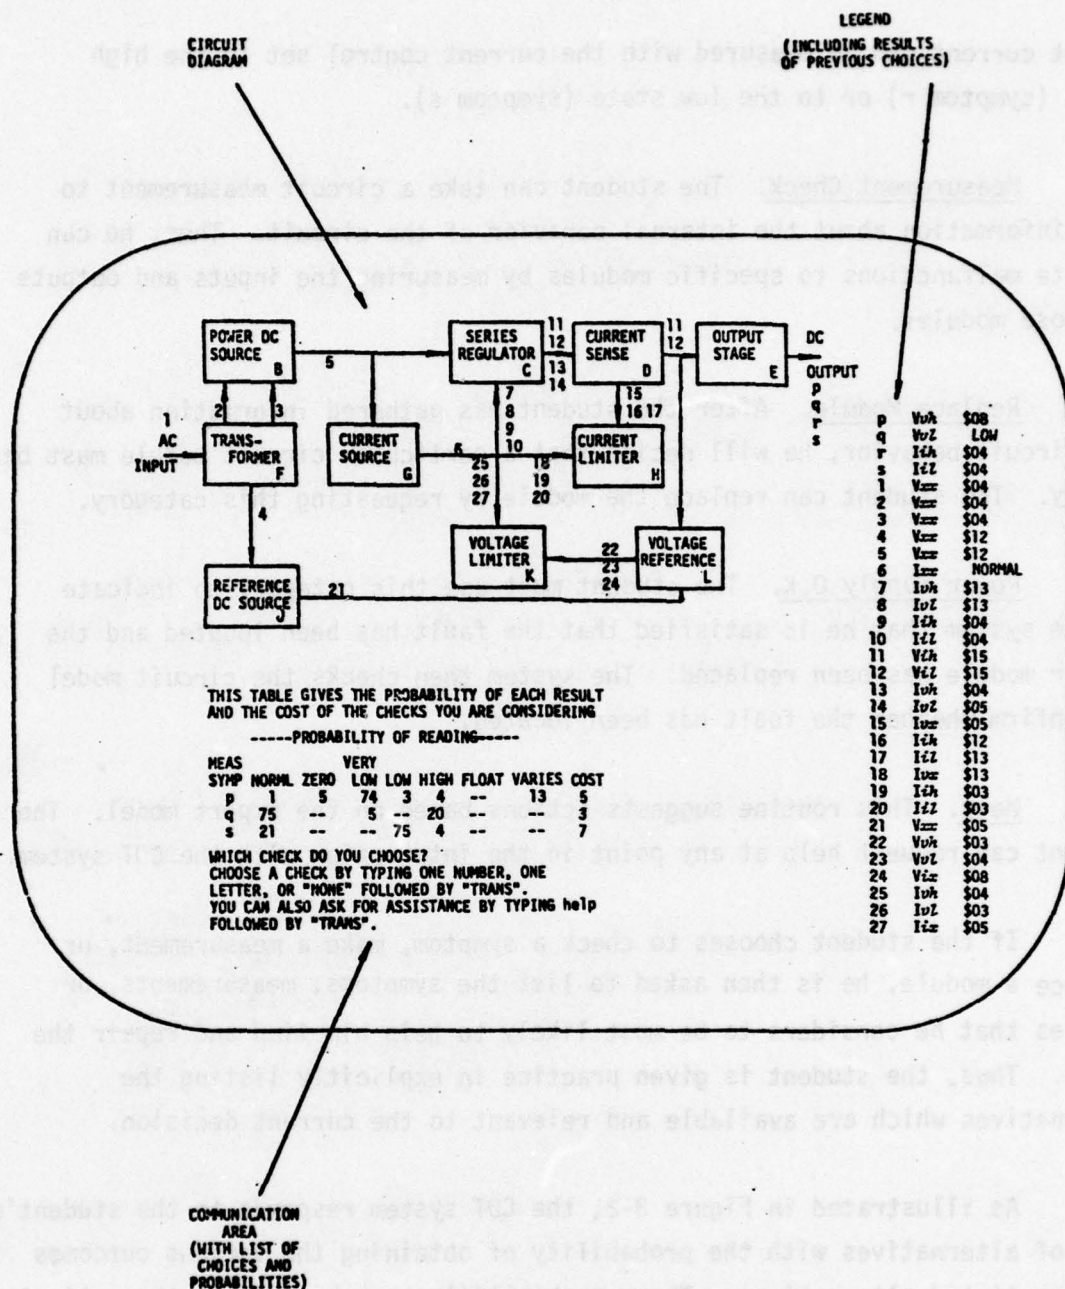
Replace Module. After the student has gathered information about the circuit behavior, he will decide that a particular circuit module must be faulty. The student can replace the module by requesting this category.

Power Supply O.K. The student must use this category to indicate to the system that he is satisfied that the fault has been located and the proper module has been replaced. The system then checks the circuit model to confirm whether the fault has been located.

Help. This routine suggests actions based on the expert model. The student can request help at any point in the interaction with the CDT system.

If the student chooses to check a symptom, make a measurement, or replace a module, he is then asked to list the symptoms, measurements, or modules that he considers to be most likely to help him find and repair the fault. Thus, the student is given practice in explicitly listing the alternatives which are available and relevant to the current decision.

As illustrated in Figure 3-2, the CDT system responds to the student's list of alternatives with the probability of obtaining the various outcomes for the listed alternatives. These probabilities are based upon the estimates of the probabilities, given what is already known about the problem. Thus the probabilities are revised as more symptom and measurement checks are made.



**STUDENT'S DISPLAY
OF A PARTIALLY-COMPLETED PROBLEM**

FIGURE 3-2

The student chooses a particular action from among the alternatives listed in the table of probabilities. After a measurement is chosen, the result of taking the measurement is shown in the legend. The results of checking symptom q and taking measurement 6 are illustrated in Figure 3-2. The legend thus serves as a "scratch pad" with a record of previous actions for the current problem.

Rather than make actual circuit measurements, the student simply chooses the action by typing on the keyboard and the task model simulator provides the result of the symptom check or circuit measurement.

In deciding on his course of action, the trainee has to weigh the cost involved in each of the choices, and the probability that a particular choice will help him find and repair the actual fault. Since costs and payoffs are the basic currency of decision making, and decisions must usually be made in the face of uncertainty, the trainee is immediately confronted with the main elements of the decision problem.

Help Routines. The "help" option is provided as an additional method of instruction. These routines provide suggested actions to direct the student toward appropriate actions. These routines are based on the expert model. The help which can be given is listed in three categories. First, the training system will list the circuit modules which are most likely to be faulty. Second, the system can suggest the action that the expert would choose. Third, the help routine will suggest which action the expert would choose from among the alternatives that the student is considering.

The CDT system can be divided into four functional parts: (1) Task Simulator; (2) Expert Model; (3) Student Model; and (4) Instructional Logic.

3.3 Task Simulator

The present task involves troubleshooting an electronic power supply; however, the following discussion could apply to a wide variety of diagnostic decision tasks (e.g., medical diagnosis, automotive repair, troubleshooting microelectronics, etc.). The task is based on a block diagram of the power supply circuit. The behavior of the power supply is simulated with a table of faults symptoms, and measurement outcomes. This method does not depend on circuit theory and thus could be generalized to subject areas for which functional equations do not exist (e.g., medical diagnosis). The task simulator generates circuit faults and simulates the results of checking symptoms, taking measurements, and replacing modules. Measurement results are presented in a semi-interpreted form (high, normal, low)(similar to the approach of Bond and Rigney, 1966), rather than as absolute readings so that the student need not refer to a table of normal circuit levels. These simplifications do not affect the inherent judgmental nature of the troubleshooting task.

The fault model is a subset of the task simulator that is used to provide updated module failure and measurement outcome probabilities. This fault model is based on a state space approach of calculating the conditional probabilities. For the purposes of the fault model, the normal condition is considered to be a fault. Two conditional probabilities are calculated:

$$P(X_{ij}|\bar{X}) \quad (3.1)$$

and $P(M_\ell|\bar{X}) \quad (3.2)$

where

X_{ij} is a measurement outcome to be obtained

\bar{X} is the vector of all measurement outcomes obtained
(measurement history)

M_ℓ is module, ℓ , containing a fault.

The conditional probability shown in equation 3.1, the probability of getting result X_{ij} , given the measurement history, is calculated as follows: Let K be an index over faults. $P(K)$ is the probability of fault K . For each fault, K , each measurement, X_j , must have a unique outcome, X_{ij} . Then let L_{ij} be the set of faults associated with measurement outcome X_{ij} .

Then

$$\sum_k P(K) = 1 \quad (3.3)$$

and the initial measurement probabilities (before any other measurements have been taken) are calculated as

$$P(X_{ij}) = \sum_{[k \in L_{ij}]} P(K) \quad (3.4)$$

Also, let L_x be the set of faults still possible, given the measurement history \bar{X} . Then

$$L_x = \bigvee_{L_{ij} | X_{ij} \in \bar{X}} \quad (3.5)$$

and the conditional measurement probability is calculated as

$$P(X_{ij} | \bar{X}) = \frac{\sum_{K \in (L_{ij} \cap L_x)} P(K)}{\sum_{k \in L_x} P(K)} \quad (3.6)$$

After a measurement result, X_{ij} , is obtained, the fault probabilities, $P(K)$, are adjusted to $P(K|X_{ij})$ as follows:

$$P(K|X_{ij}) = 0 \quad (3.7)$$

for all $K \notin L_{ij}$ and

$$P(K|X_{ij}) = \frac{P(K)}{\sum_{[K \in L_{ij}]} P(K)} \quad (3.8)$$

The probability of the module failure is then calculated

$$P(M_x) = \sum_{K \in L_x} P(K) \quad (3.9)$$

The conditional probabilities, $P(X_{ij}|\bar{X})$ and $P(M_x|\bar{X})$, are used by the expert model in calculating the expected value for each measurement or module replacement. These probabilities are also displayed to the student during his process of considering subsets of measurements or modules.

The simulated task may be altered by specifying a new set of task problems and the associated probabilities. The simulation also requires that the set of possible task actions be specified and permits a set of costs for these actions to be included. This state representation technique contrasts with a task simulation based on a set of defining equations. The former technique permits simulation of a wide range of tasks, including those for which functional equations do not exist. However, the latter technique is powerful in terms of the dynamic range of simulated task characteristics.

3.4 Expert Model

The Expert Model is a fixed parameter model with the following capabilities:

- (a) It can select the decisions with the highest expected utility.
- (b) It can be used to suggest an action from among those being considered.
- (c) Its utilities can be used as a standard to which the student's utilities can be compared.

An expected utility (EU) decision model is used to represent the expert. The EU is a prescriptive model which makes use of a criterion for optimum choice among alternatives, assuming "rational behavior" (Edwards, 1962; Fishburn, 1964). The choice criterion employed is the maximization of individual expected relative utilities as obtained by a weighted sum of individual utilities of consequences and their probability of occurrence.

The maximum EU principle has become a widely acceptable normative decision model for risky decision making (Luce and Raiffa, 1957; Krantz et al, 1971). The work of Tversky (1967); Goodman et al (1967); and others has indicated that the expert maximization principle provides a good first approximation for decision making under risk.

More specifically, the expected utility of an action is

$$EU_j = \sum_i P_{ij} U_{ij} \quad (3.10)$$

P_{ij} = probability that the i^{th} consequence in a set of n consequences will occur if action A_j is selected by the decision maker.

U_{ij} = relative utility of the i^{th} consequence of the j^{th} action.

The EU model is used here as a basis for defining optimum strategies and as a structure for adaptively estimating the student's utilities as inferred from his decision behavior. Using an on-line adjustment of the specific student utility structure, an adaptive learning network searches for the student utilities which can explain his action by the criterion of maximization of EU. Thus, the model continuously tracks the student's decision strategy as it changes during the course of training. In the instructor model, a fixed EU model is used in real time as a criterion for recommending actions to the student in response to the student's requests for assistance. Off-line, an adaptive EU model is used to estimate the instructor's values.

3.4.1 Information Gain. The expected value model itself is insufficient to model dynamic decision behavior where the primary goal of the task is more information gain. Such is the case with tasks of diagnosis, fault detection, intelligence gathering, troubleshooting, etc. In these instances, the expected value of an action is also a function of the information to be gained if the action is selected. Shannon and Weaver's (1949) measure of information,

$$I = -\sum_k P(K) \log_2 P(K) \quad (3.11)$$

is a commonly used measure of information.

We need to find an information gain function, α_{ij} , which associates an information gain with a measurement outcome, X_{ij} . This function is incorporated into the EV equation as follows:

$$EV = \sum_i \alpha_{ij} P_{ij} V_{ij} P_{ij}(\alpha_{ij} U_{ij}) \quad (3.12)$$

Since we are modeling a subjective decision process, there are several possible forms of the information gain function, α_{ij} . Some of these are discussed below:

$$(1) \alpha_{ij} = \sum_k [P(K|X_{ij}) \log_2 P(K|X_{ij}) - P(K) \log_2 P(K)] \quad (3.13)$$

With this calculation of α_{ij} , the value for an outcome is weighted by the information to be gained about remaining faults if the outcome is obtained.

(2) A similar function to 3.13 is calculated by substituting $P(M_k)$ for $P(K)$. In this instance, the value for the outcome is weighted by information to be gained about faulty modules.

(3) A direct substitution of Shannon and Weaver's measure of entropy gives

$$\alpha_{ij} = -\sum_i P_{ij} \log_2 P_{ij} \quad (3.14)$$

However, the expected value of an action becomes

$$\begin{aligned} EV_j &= \sum_i P_{ij} (-P_{ij} \log_2 P_{ij}) V_{ij} \\ &= \sum_i -P_{ij}^2 \log_2 P_{ij} V_{ij} \end{aligned} \quad (3.15)$$

which multiplies the weighting of the function by the outcome probabilities. Hence a better formula for the information gain is

$$\alpha_{ij} = \sum_i -\log_2 P_{ij} \quad (3.16)$$

$$\text{and } EV_j = \sum_i P_{ij} (-\log_2 P_{ij}) V_{ij} \quad (3.17)$$

Note that equation 3.16 has the same value for all outcomes, i , of measurement, j . In this case knowledge of how outcomes affect fault probabilities are not needed.

$$(4) \alpha_{ij} = 1 \quad (3.18)$$

is valid where the specific task does not involve information acquisition. It sometimes represents the behavior of beginning troubleshooters who tend to prefer to verify existing knowledge.

- (5) Another approach is taken by the troubleshooter who tries to eliminate the largest number of possible faults (regardless of probability or module in which the fault occurs) with each measurement. The following α models such a troubleshooter:

$$\alpha_{ij} = -f_i \log_2 f_i \quad (3.19)$$

where, considering previous measurements

$$f_i = \frac{\text{Faults associated with } i}{\text{All possible faults}} \quad (3.20)$$

In the initial trials with the system, the information gain function represented by equation (5) was used. The other information gain functions are also available at the experimenter's option.

3.4.2 Simulated Student. The expert model can be used to simulate a student by substituting student values for expert values and using an option in the program which automatically goes through the troubleshooting process as a student would. The simulated student is used for the following purposes:

- (1) As a debugging aid.
- (2) To simulate proposed experimental sequences.
- (3) To compare actual student behavior with model behavior to test model validity.

Any set of values (e.g., values estimated from a human subject's performance, or theoretically ideal values) can be used to form the basis of the decision behavior which the program generates. This subprogram is an excellent aid for debugging software, and for quickly finding possible problems with a proposed experimental sequence before real subjects are used. The simulated student subprogram is also a useful analysis tool since it can be used to generate consistent behavior so that system parameters can be fine tuned or to serve as an optimized consistent standard to which student performance can be compared. It has the future potential of demonstrating a particular type of decision behavior to the student.

3.3.4 Requests for Help. A "help" routine is available which uses the ideal response characteristics of the instructor EU model. The algorithms interpret a student's request for assistance, and responds to student inquiries by performing the following functions:

- (1) Listing potentially faulty modules.
- (2) Suggesting an optimum action, irrespective of whether the student has considered it.
- (3) Suggesting an optimum action from the set of actions the student is already considering.

Function 1 involves checking the current action-outcome probabilities to deduce which modules have a significant probability of being faulty at the present state in the debug cycle. An optimum action is chosen in Function 2 by selecting that action with the maximum expected utility. The expected utilities of all actions are calculated from the expert's utilities and the probabilities for action outcomes. Function 3 is similar to function 2; however, the action with the maximum expected utility is selected from the subset of actions that the student has selected for consideration.

3.5 Student Model and Dynamic Utility Assessment

The student model has the same mathematical form as the expert model except that the values, V_{ij} , are adaptively determined by the methods of dynamic utility assessment. The EV equation is used to predict the student's action.

The student model is designed to model subjective aspects of the student's decision behavior by continuously "tracking" changes in the student's values in real time. These estimated values provide the basis for instructing the student about appropriate evaluation of action outcomes. During the process of converging upon stable estimates of the student's values, the adaptive student model can also provide early estimates of widely inappropriate values. Immediate instructions regarding these values can direct the student toward more consistent decision-making while he becomes accustomed to the training problems.

Since the student must also learn to focus on the odds of action outcomes, the training system displays the probabilities to the student. The student model uses the same set of probabilities as used in the expert model.

The dynamic value estimation technique, developed by Perceptronics in the context of a decision aiding task (Freedy, Weisbrod, Davis, May, and Weltman, 1974), is based on the principle of a trainable multi-category pattern classifier. The value estimator observes the operator's choices among R possible decision options available to him, viewing his decision making as a process of classifying patterns by means of an expected value evaluation, or discriminant, function. These classifications are compared with the operator's decisions and an adaptive error-correction training algorithm is used to adjust pattern weights, which correspond to values, whenever the classifications are incorrect. Thus, the value estimator "tracks" the operator's decision making and "learns" his values.

A multi-category pattern classifier (Nilsson, 1965) receives patterns of data and responds with a decision to classify each of the patterns in one

of R categories. The classification is made on the basis of R linear discriminant (or evaluation) functions, each of which corresponds to one of the R categories. The discriminant functions are of the form

$$g_i(\bar{X}) = \bar{W}_i \cdot \bar{X} \text{ for } i = 1, 2, \dots, R \quad (3.21)$$

where \bar{X} is the pattern vector and \bar{W}_i is a weight vector. The pattern classifier computes the value of each discriminant function and selects the category, i , such that

$$g_i(\bar{X}) > g_j(\bar{X}) \quad (3.22)$$

for all $j = 1, 2, \dots, R; i \neq j$.

The adaptive error-correction training algorithm is very straightforward. Whenever the category selected by the pattern classifier, i , is different from the actual classification, k , the weights \bar{W}_i are adjusted to reduce (punish) the value of $g_i(\bar{X})$ and the weights \bar{W}_k are adjusted to increase (reward) the value of $g_k(\bar{X})$. Thus,

$$\bar{W}_i' = \bar{W}_i + \gamma \cdot \bar{X} \quad (\text{Reward}) \quad (3.23)$$

$$\bar{W}_i' = \bar{W}_i - \gamma \cdot \bar{X} \quad (\text{Punish}) \quad (3.24)$$

where γ is the correction increment.

The dynamic value estimator, schematically represented in Figure 3-3 classifies pattern vectors

$$\bar{P} = [p_{1,1}, p_{1,2}, \dots, p_{i,j}] \quad (3.25)$$

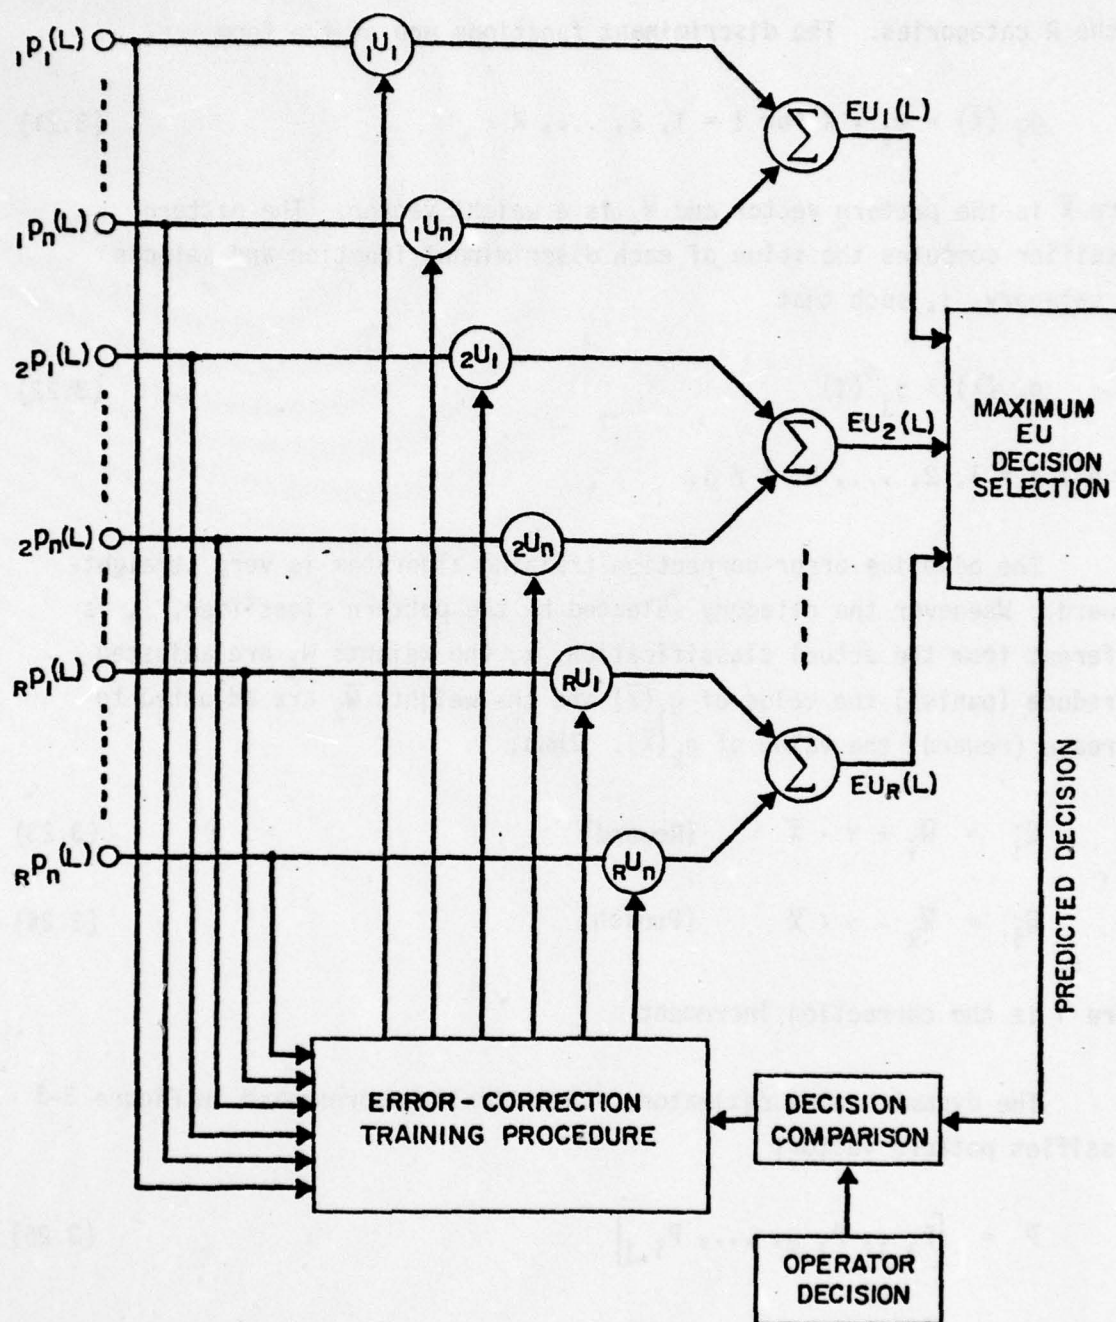


FIGURE 3-3. SCHEMATIC REPRESENTATION OF DYNAMIC VALUE ESTIMATOR

whose components, p_{ij} , are the aggregated probabilities of the result, i , of action j . The discriminant functions are the expected values

$$EV_j = \sum_i p_{ij} \cdot V_{ij} \alpha_{ij} \quad (3.26)$$

of the actions. The value estimator computes the EV of each action and selects that action for which the EV is maximum.

A fixed increment training rule is used to adjust the values. Whenever the action, d , selected by the value estimator differs from the action, c , selected by the student, the values associated with the estimator action are punished and those associated with the student are rewarded:

$$V_{jd}^{t+1} = V_{jd}^t - (\gamma \cdot p_{jd}) \alpha_{ij} \quad (\text{Punish}) \quad (3.27)$$

$$V_{jc}^{t+1} = V_{jc}^t + (\gamma \cdot p_{jc}) \alpha_{ij} \quad (\text{Reward}) \quad (3.28)$$

The values at time $t+1$ are computed for all action results, j . The correction increment, γ , is a constant which can be adjusted to give optimum convergence of the estimated values.

3.6 Instructional Logic

Comparison of parameters in the student model with those in the expert model leads to the instructions and feedback given to the student. The instructional logic compares the utilities of the expert and the student models to generate instructional feedback in the form of generating fault problems and offering instructional "help" to the student. Help is offered by monitoring the student's decisions and using the instructor model to give prompts to the trainee.

The extraction of prescriptive information from the student's utilities is the central focus in the CDT system concept. In particular, the analysis of utilities provides a direct measure of student's decision-making consistency, an indication of whether he approaches the correct utilities, and a measure of the rates at which he approaches the correct utilities. Utility discrepancies between that of the student and the instructor model provide direct diagnostic information regarding student decision strategies and areas where special instruction are required.

An analysis of the expert technician's performance identified the critical measurements for efficient diagnosis and fault isolation. The estimated utilities for the normal outcomes of these critical measurements are characterized by a rapid divergence from the initial starting values. Additionally, the measurement utilities show early movement toward final relative ranking. Thus, the derivative of the estimated utilities, $\frac{\delta u}{\delta t}$, can be used as an effective index for selecting instructions. During the initial selection and evaluation of instructions, statements printed on index cards were presented to the student. The experimenter selected the cards on the basis of stated rules and calculations from the printed values.

The new CDT system will contain software mechanisms for evaluating student behavior and giving the student corrective feedback. Behavior is monitored by evaluating "key" utilities in the student model for relative rank after a small number of problems. Specific differences between the student's rank and "optimum" rank are associated with specific instructions designed to "correct" the relative utilities the student places on the measurements involved.

The procedure of having the student perform some problems and subsequently evaluating his utilities and providing feedback for him is repeated for as many instructional phases as the operator wishes to specify. The key measurements, instructions available, and instructional criteria is phase specific and may be set by the operator.

A criterion for selecting an instruction is specified as a Boolean expression whose truth value is determined on the basis of most recent student behavior. If the truth value of the expression is FALSE, the corresponding instruction is not given. If the value is TRUE this instruction is given.

As an example of the instruction-selection procedure, consider the following situation: suppose that a particular instruction is to be given if and only if the utility for measurement 4 is higher than the utilities for measurements 6 or 7, and the utility for measurement 3 is NOT less than the utility for measurement 6. The corresponding Boolean expression is $((4P7) \vee (4P6)) \wedge \neg(6P3)$. Assume that 4 is indeed greater than 7, 4 is greater than 6, and that 6 is greater than 3. In this example, the instruction would not be given since the operation $\neg(6P3)$ was evaluated as FALSE.

Currently the elements of the Boolean expression are precedence relations for measurement utilities. The Boolean interpreter will, however, evaluate additional elements with minor additions to the system so that additional instructional criteria, based on utilities and/or on other parameters, may be developed as a basis for more sophisticated instructional feedback.

Currently, the derivative of the estimated student utilities is not being used as a basis for instructional feedback, this may be added later. The ranking of utilities, after a few problems are performed by the student, provides more information than the direction of change of utilities and is currently used instead. Under certain circumstances (such as when a pre-determined series of problems is given to the student), valid directional information from utilities may precede valid rank information and utility derivative criteria would be justified on the basis of more rapid responsiveness to student behavior. If so, such criteria and associated instructions will be added to the system.

The interpreter currently processes 4 Boolean operators and one relation. These are:

Operators

V - OR
Λ - AND
X - EXCLUSIVE or
- - Not

Relations

P - precedes (ranked higher)

Other relations and operators may be added easily to the system.

3.7 Diagnostic Report of Student Performance

To fully evaluate student and model performance it was necessary to develop a data collection package, a set of performance measures, and a student diagnostic report. The data collection package was designed to provide a general purpose tool for collecting experimental data in a form which is amenable to printing, storing on disk, or using for computer analysis. The program uses the concept of inserting data collection points in the flow of control of an experimental system. At each data collection point the values of specified variables are collected. Each data collection point may have associated with it a heading which is printed above the data for the corresponding data collection point. Also, each variable has an optional 16 character label associated with it which is printed in the left most part of the page.

In addition to the summary of activities provided by the data collection routine, several performance indices were developed to characterize the student's performance during the interaction with the CDT system. These performance measures are printed at the conclusion of each problem along with the problem summary data. All measures are calculated as mean values across each fault problem.

The performance measures include:

- (1) Relative Competence -- A measure comparing the expected utilities of the student and expert for the chosen action. Competency is given by the following equation:

$$\text{COMPETENCE} = \frac{\text{EU}^e (\text{student choice})}{\text{EU}^e (\text{model choice})} \times 100$$

where EU^e is an expected utility calculated from the utilities of the expert model.

- (2) Relative Consistency -- A measure of the expected utility of the student's choice expressed as a percentage of the expected utility of the model prediction. This compares the student's own performance with the performance of his adaptive model.

$$\text{CONSISTENCY} = \frac{\text{EU}^s (\text{student choice})}{\text{EU}^s (\text{model choice})} \times 100$$

Consistency is a measure of the stability of the decision maker's behavioral patterns. Untrained DM's often alter their tactics using inappropriate criteria. Often they modify their behavior on the basis of inadequate number of preceding trials (Sidorsky & Simoneau, 1970). Thus, decision consistency is an important measure which defines whether the DM has reached a stable tactical decision strategy.

The adaptive decision modeling system inherently provides a measure of decision consistency. If the DM's behavior is inconsistent, the model-derived estimates of his utilities will not converge, whereas stable behavior will lead to convergence. Thus, the degree to which the utilities converge also provides a direct measure of decision consistency.

- (3) **Relative Information Gain** -- Indicates the amount of information the subject obtained compared to the maximum information available.

$$\text{ALPHA} = \frac{\text{Information Gain of Student's Choice}}{\text{Information Gain of Choice with Maximum Information}} \times 100$$

- (4) **Relative Information Gain of Considered Alternatives** -- Indicates the amount of information in the several actions that the student considers before each decision, as compared with the information potentially available from the set of actions with the maximum expected values.

$$\text{ALPHA (CONS)} = \frac{\frac{\sum (\text{Information Gain of Considered Actions})}{n}}{\frac{\sum (\text{Information Gain of Maximum EU Actions})}{n}} \times 100$$

Table 3-1 gives a list of additional performance summary indices that are printed by the data collection routines. These indices are printed at the completion of each problem.

3.8 System Software

The CDT system software is written in assembly language and contains a number of modules. The software is briefly described in the first annual report and is fully documented in a programmer's reference manual.

Information Structures. The state of the instructional sequence at any given time is maintained in information structures containing the following information:

TABLE 3-1. PERFORMANCE SUMMARY INDICES

1. FIRST LINE: Information typed by experimenter at beginning of experiment (e.g., Name, Date, Session, etc.).
2. PROBLEM: Problem number.
3. FAULT: Fault number.
4. BAD MODULE: Faulted module.
5. PROBLEM COST: Cost to complete problem.
6. TOTAL COST: Cost to complete all problems so far.
7. DECISIONS: Number of decisions taken to complete problem.
8. ADJUSTMENTS: Number of times student model failed to predict during current problem.
9. COMPET.: Competency of student's decisions.
10. CONSIS.: Consistency of student's decisions.
11. ALPHA: Effective information gain of student's measurement choices.
12. ALPHA (CONS): Information gain of student's considerations.
13. TIME (MIN): Time (in minutes) to complete current problem.
14. NO. TIMES NONE USED: Number of times student used "NONE" in current problem.
15. NO. TIMES HELP USED: 1, 2, 3: Number of times each of the three help options were used in this problem.

- . The actual circuit fault.
- . The measurement results obtained to this point.
- . Action phase (taking a measurement, replacing a module, help, declare operational, checking a symptom).
- . Total expended cost.
- . Actions currently being considered, if any.
- . Actions currently chosen, if any.
- . Student's current values for measurement results and module replacements, as represented in the student model.

Fundamental to the system concept is the fact that the state of the instructional sequence can be represented by variables taking on discrete values. For this reason the state can be very compactly represented (a fact that would be very useful in a multi-student environment). The measurement results so far obtained, the actions considered, and the actions chosen are each represented by binary vectors.

3.9 Instructor/Computer Interactions

The instructor interacts with the CAI system primarily through the teletype. Using an interactive control program the instructor may easily modify the nature and complexity of the task environment, the decision model performance characteristics, and the structure of the student/computer interface. Additionally, he may save the student environment in the middle of a testing session and restore it at a later time to resume testing.

The instructor controls the task environment by modifying the characteristics of the measurements and the circuit faults. The experimenter can modify the fault behavior of the circuit by changing the probability values. He can add new faults and measurements by making additional entries in the tables which define the information structures.

The performance characteristics of the decision model can be controlled by modifying (1) the initial value levels used by the adaptive EU model, (2) the learning rate of the utility estimator, and (3) the EU evaluation function used by both the model and the utility estimator. The easiest to modify are the initial settings of the value matrix, which are input by the instructor during program initialization. These initial values affect the behavior of the adaptive EU model and the utility estimator, at least during the early stages of a run. The learning rate of the utility estimator is controlled by a correction increment. This parameter affects the rate of convergence of the utility estimator and determines its sensitivity to changes in the operator's decision behavior. The size of the correction increment also affects the amount of variance which will result from inconsistent operator behavior. A more difficult method of controlling the decision model is modification of the expected value function. This function, also used as a discriminant function by the utility estimator, is programmed into the system. Such modification of the EU function might be done, for example, if new types of task actions were required in the training system.

The interactive control program allows the experimenter to alter or select a large number of system parameters. This program is designed to allow additional commands to be added easily when required. Basically, three types of function are supported: (1) setting, resetting, changing, or displaying current CDT system parameters; (2) saving the current problem environment; and (3) recording the student's behavioral history or rerunning a previously saved history to exactly reproduce a student's behavior. In the first category, some of the implemented options are: (1) suppressing the use of "NONE" and "HELP" by the student; (2) assigning of a cost to "NONE" and "HELP"; (3) changing of the cost of symptoms, measurements, or modules; (4) choosing an information gain function; (5) selecting the order in which faults are chosen (if not random); (6) setting the number of measurements the student may consider at one time; (7) running a "simulated student"; (8) resetting the student's or the expert's values; and (9) suppressing probability

presentation. All parameters can also be displayed, punched on paper tape (or to a disc file), or read from paper tape, as well as being changed at any time, or cause it to be invoked after a certain number of problems have been completed.

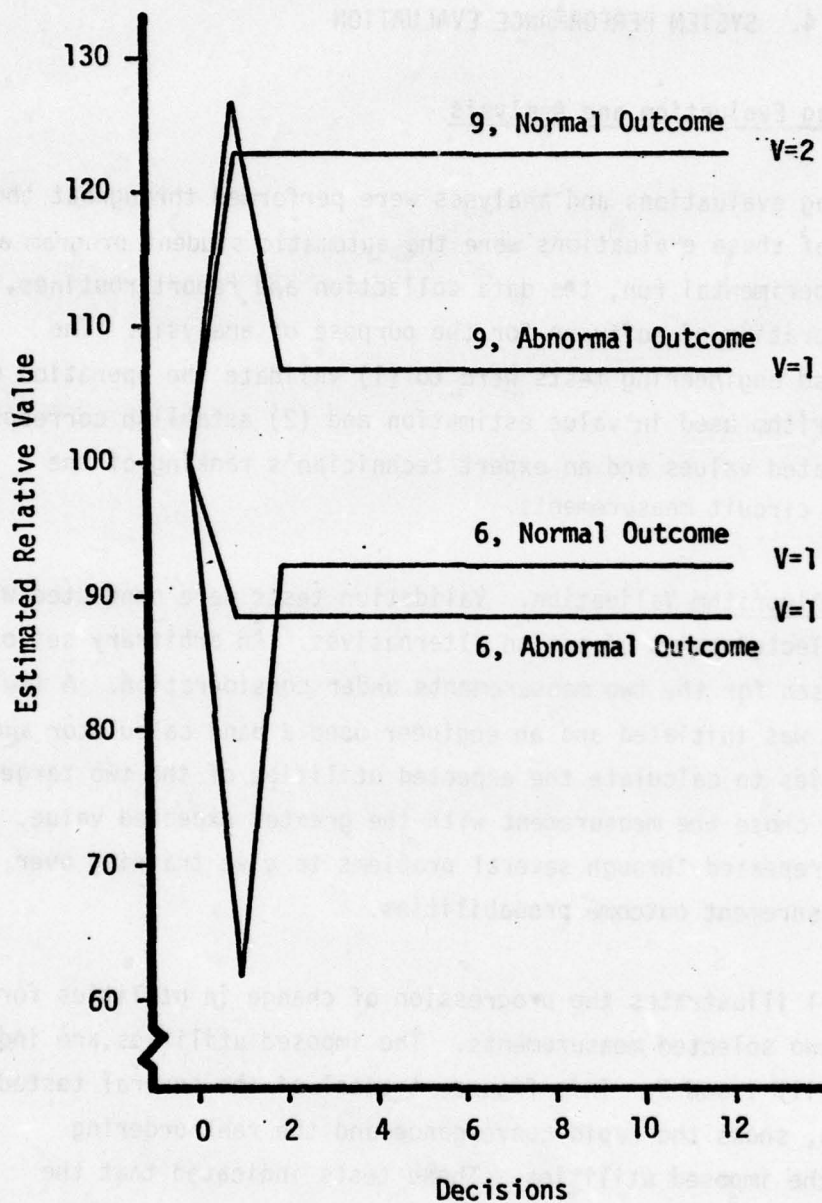
4. SYSTEM PERFORMANCE EVALUATION

4.1 Engineering Evaluation and Analysis

Engineering evaluations and analyses were performed throughout the year. The tools of these evaluations were the automatic student program as used to simulate an experimental run, the data collection and report routines, and the temporary alteration of software for the purpose of analysis. The objectives of these engineering tests were to (1) validate the operation of the training algorithm used in value estimation and (2) establish correspondence between the estimated values and an expert technician's ranking of the importance of the circuit measurements.

4.1.1 Training Algorithm Validation. Validation tests were conducted with local tests of selected pairs of action alternatives. An arbitrary set of utilities was chosen for the two measurements under consideration. A fault-isolation problem was initiated and an engineer used a hand calculator and his chosen utilities to calculate the expected utilities of the two target measurements. He chose the measurement with the greater expected value. This process was repeated through several problems to give training over a wide range of measurement outcome probabilities.

Figure 4-1 illustrates the progression of change in utilities for the two outcomes of two selected measurements. The imposed utilities are indicated as being arbitrarily 1 and 2. This figure, typical of the several tested measurement pairs, shows the rapid convergence and the rank ordering consistent with the imposed utilities. These tests indicated that the utility training algorithm was operating correctly for isolated values. A global test was designed to analyze the algorithm function in the complete measurement set.



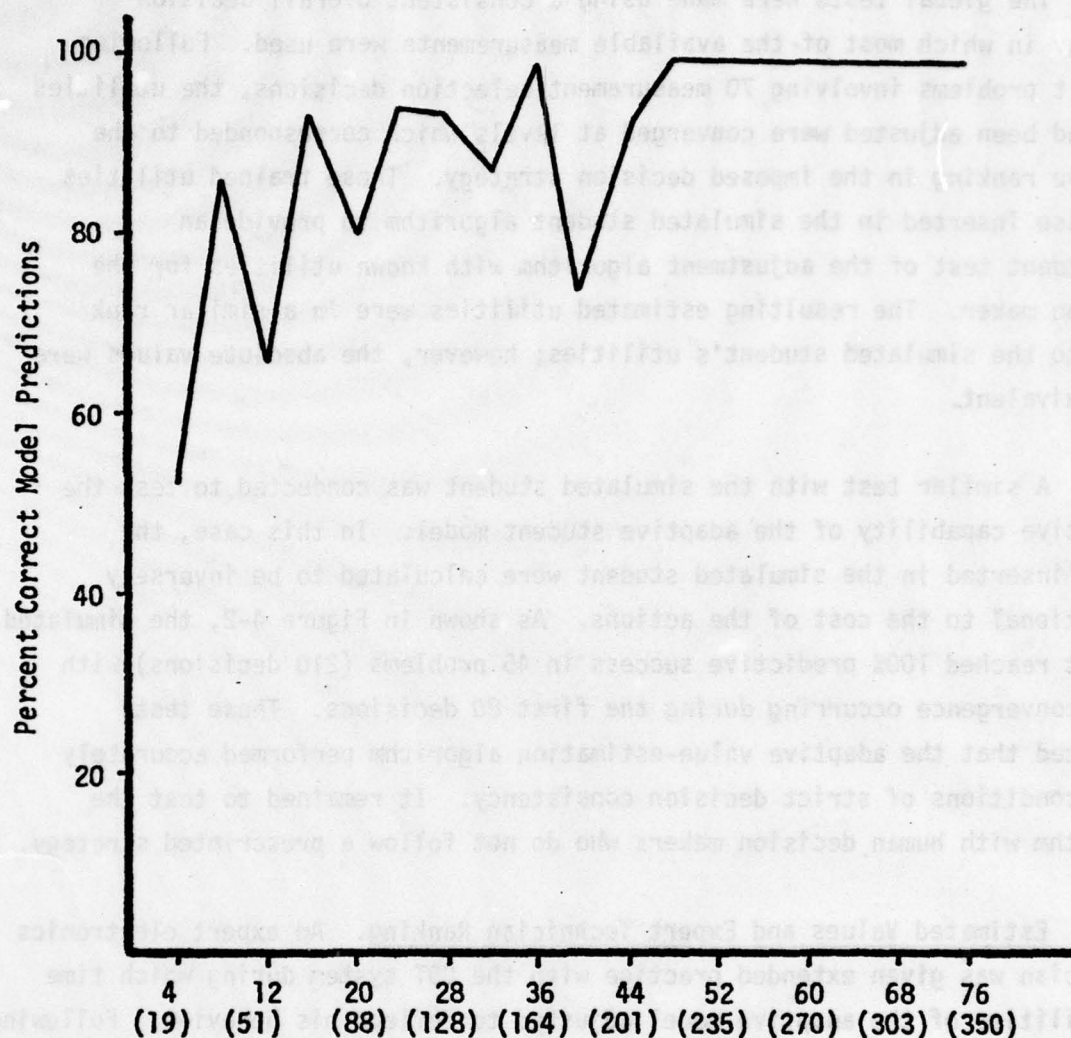
ESTIMATED VALUES FOR TARGET
MEASUREMENT OUTCOMES
(V - Imposed Values)

Figure 4-1

The global tests were made using a consistent overall decision strategy in which most of the available measurements were used. Following 14 fault problems involving 70 measurement-selection decisions, the utilities that had been adjusted were converged at levels which corresponded to the relative ranking in the imposed decision strategy. These trained utilities were also inserted in the simulated student algorithm to provide an independent test of the adjustment algorithm with known utilities for the decision maker. The resulting estimated utilities were in a similar rank order to the simulated student's utilities; however, the absolute values were not equivalent.

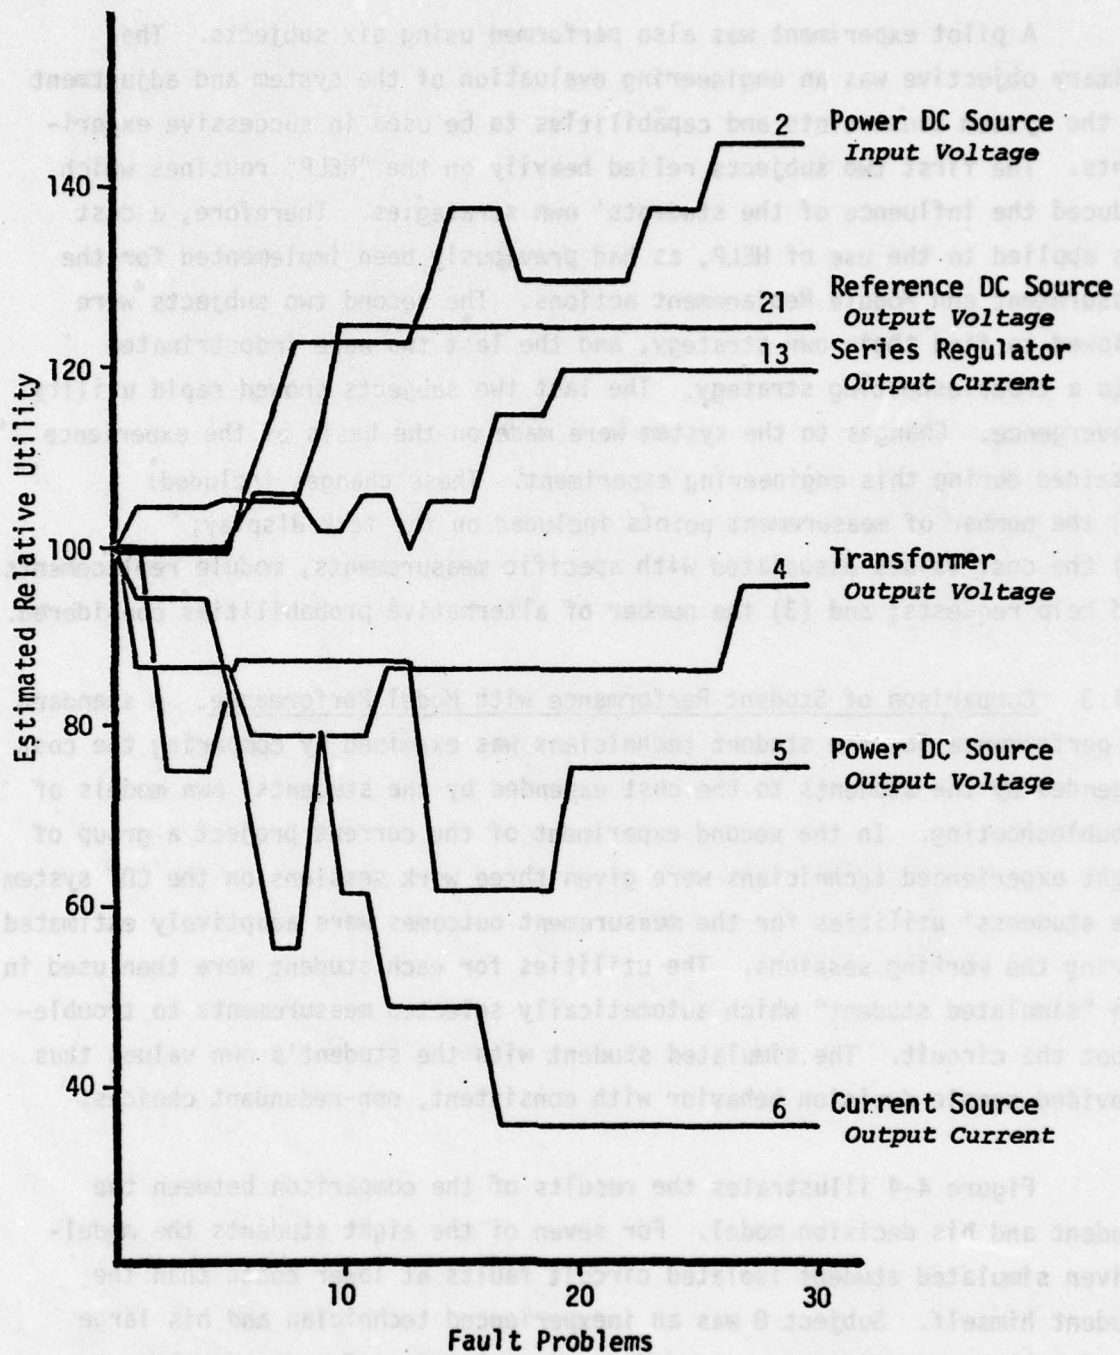
A similar test with the simulated student was conducted to test the predictive capability of the adaptive student model. In this case, the values inserted in the simulated student were calculated to be inversely proportional to the cost of the actions. As shown in Figure 4-2, the simulated student reached 100% predictive success in 45 problems (210 decisions) with rapid convergence occurring during the first 80 decisions. These tests indicated that the adaptive value-estimation algorithm performed accurately under conditions of strict decision consistency. It remained to test the algorithm with human decision makers who do not follow a prescribed strategy.

4.1.2 Estimated Values and Expert Technician Ranking. An expert electronics technician was given extended practice with the CDT system during which time the utilities of the adaptive model adjusted to reflect his behavior. Following the extended problem solving sessions, the technician was interviewed to determine his stated reasons for selecting specific measurements and his estimates of the importance of these actions. Those measurements which he indicated had critical importance in fault-isolation were identified as the key utilities. Figure 4-3 illustrates the adjustments and convergence of utilities of the expert technician for normal measurement outcomes for key measures. These key actions were identified during the post-training interview with the expert.



**PREDICTIVE SUCCESS OF DECISION MODEL
DURING SIMULATED STUDENT PERFORMANCE**

Figure 4-2



RELATIVE UTILITY FOR KEY MEASUREMENTS
AS A FUNCTION OF FAULT PROBLEMS
(Expert Technician)

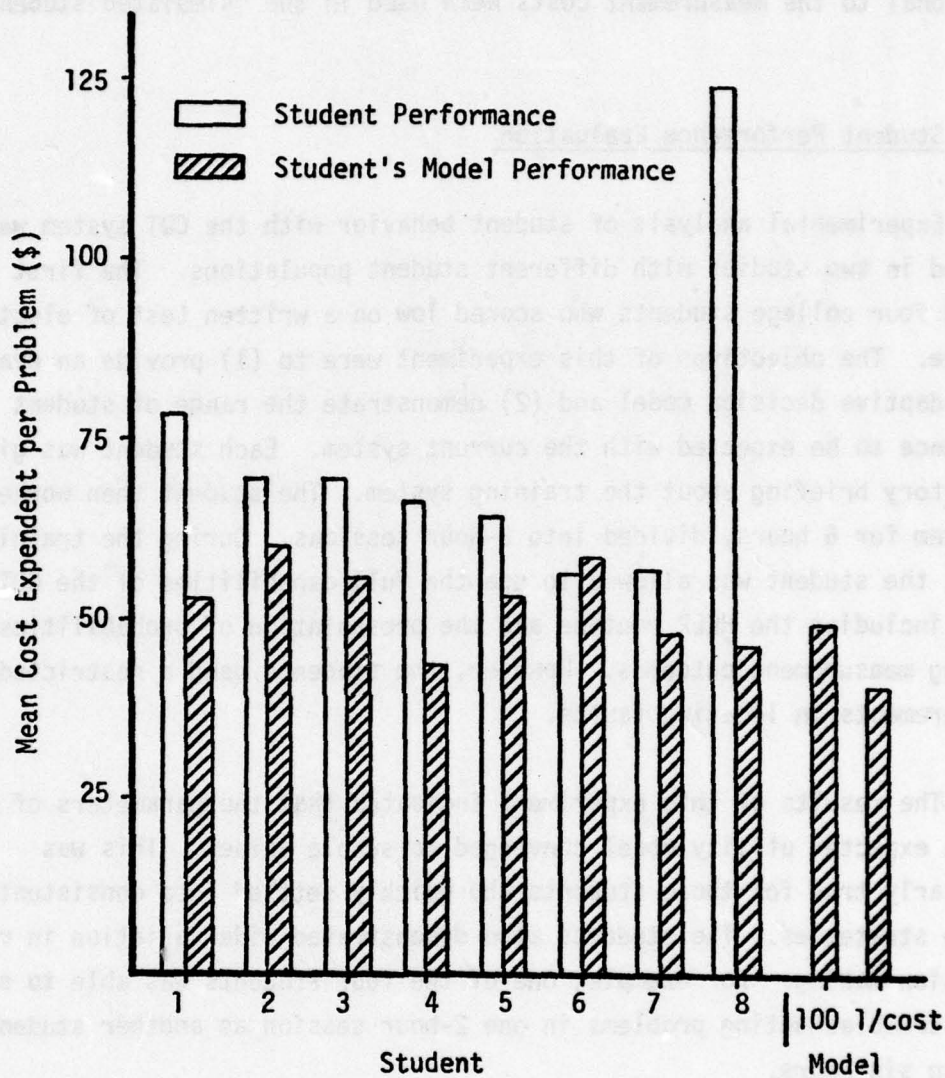
Figure 4-3

A pilot experiment was also performed using six subjects. The primary objective was an engineering evaluation of the system and adjustment of the system constraints and capabilities to be used in successive experiments. The first two subjects relied heavily on the "HELP" routines which reduced the influence of the students' own strategies. Therefore, a cost was applied to the use of HELP, as had previously been implemented for the Measurement and Module Replacement actions. The second two subjects were allowed to find their own strategy, and the last two were indoctrinated into a troubleshooting strategy. The last two subjects showed rapid utility convergence. Changes to the system were made on the basis of the experience obtained during this engineering experiment. These changes included:

- (1) the number of measurement points included on the task display;
- (2) the cost values associated with specific measurements, module replacements, and help requests; and
- (3) the number of alternative probabilities considered.

4.1.3 Comparison of Student Performance with Model Performance. A standard of performance for the student technicians was examined by comparing the cost expended by the students to the cost expended by the students' own models of troubleshooting. In the second experiment of the current project a group of eight experienced technicians were given three work sessions on the CDT system. The students' utilities for the measurement outcomes were adaptively estimated during the working sessions. The utilities for each student were then used in the "simulated student" which automatically selected measurements to troubleshoot the circuit. The simulated student with the student's own values thus provided sample decision behavior with consistent, non-redundant choices.

Figure 4-4 illustrates the results of the comparison between the student and his decision model. For seven of the eight students the model-driven simulated student isolated circuit faults at lower costs than the student himself. Subject 8 was an inexperienced technician and his large cost/problem expenditure resulted from his selection of more measurements than necessary for problem solution. For points of reference, uniform



Student and Student-Model Cost
Expenditures for Problem Solution

Figure 4-4

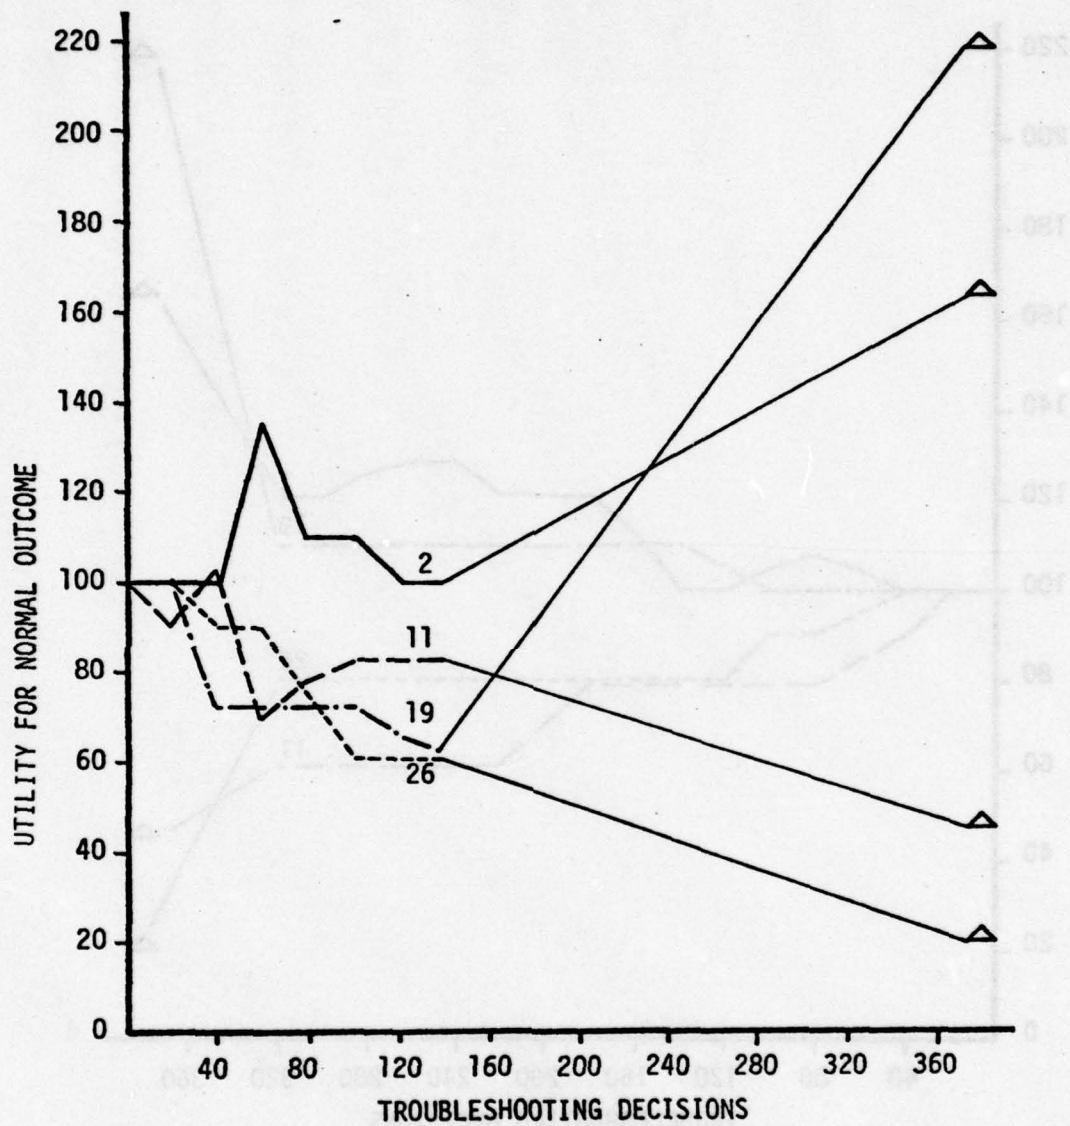
utilities for all measurement outcomes and utilities which are inversely proportional to the measurement costs were used in the "simulated student" program.

4.2 Student Performance Evaluation

Experimental analysis of student behavior with the CDT system was conducted in two studies with different student populations. The first study included four college students who scored low on a written test of electronics knowledge. The objectives of this experiment were to (1) provide an evaluation of the adaptive decision model and (2) demonstrate the range of student performance to be expected with the current system. Each student was given an introductory briefing about the training system. The student then worked with the system for 6 hours, divided into 2-hour sessions. During the training sessions the student was allowed to use the full capabilities of the CDT system, including the HELP routine and the presentation of probabilities of obtaining measurement outcomes. However, the students used a restricted set of measurements in locating faults.

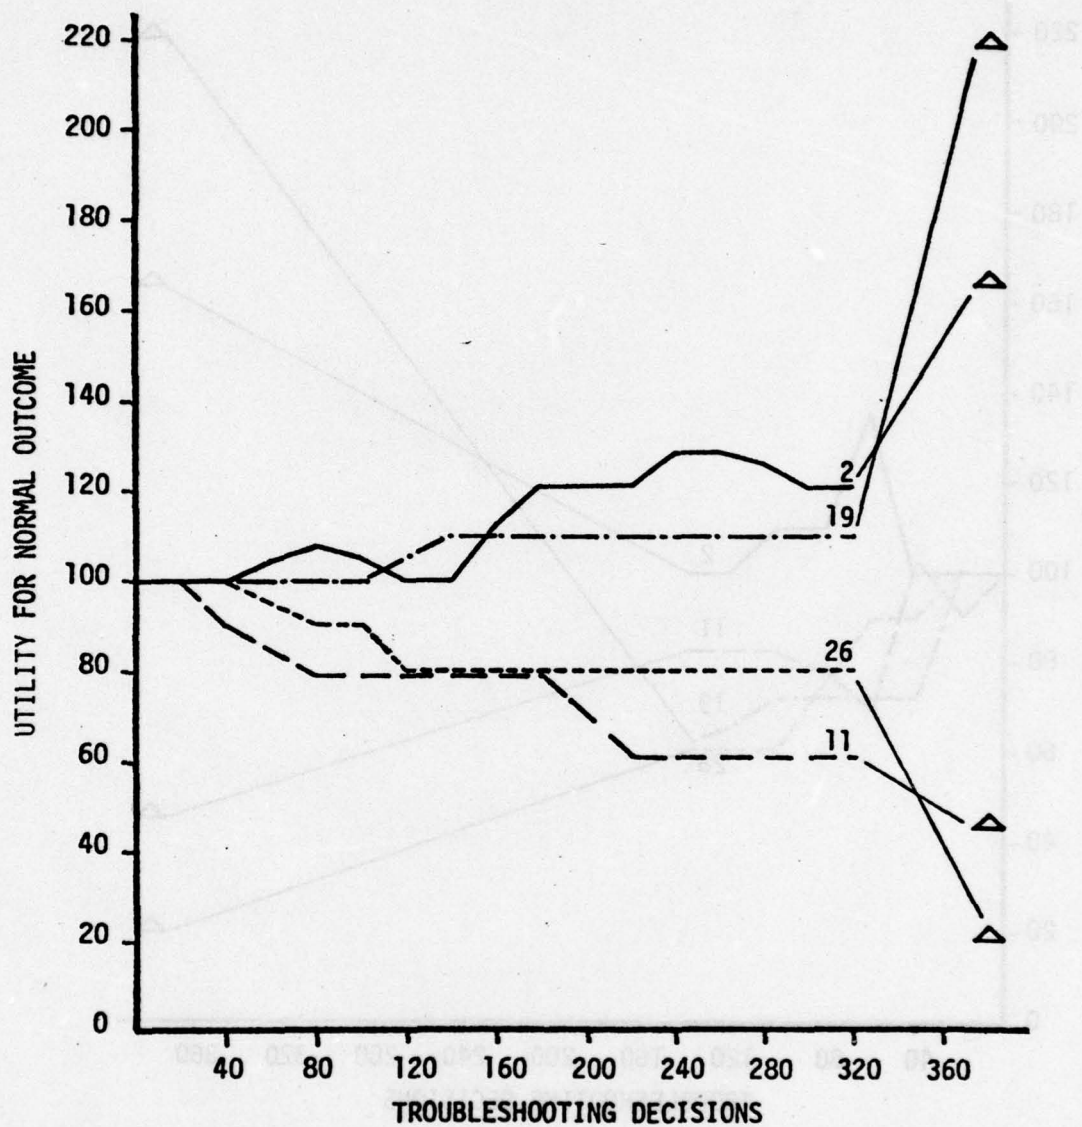
The results of this experiment indicated that the parameters of the adaptive expected utility model converged to stable values. This was particularly true for those students who quickly settled into consistent decision strategies. The students also demonstrated wide variation in rate of decision-making. For example, one of the four students was able to solve as many troubleshooting problems in one 2-hour session as another student solved in six hours.

Figures 4-5 and 4-6 illustrate the performance of the adaptive decision model for the two students described above. Utilities for several selected circuit measurements are plotted as a function of the number of student decisions. The values were set at an initial level of 100 and were adaptively adjusted to track the student's choices. For reference, values



ESTIMATED UTILITIES AS A FUNCTION OF
TROUBLESHOOTING DECISIONS
(STUDENT: A)

Figure 4-5



ESTIMATED UTILITIES AS A FUNCTION OF
TROUBLESHOOTING DECISIONS
(STUDENT: C)

Figure 4-6

calculated to be proportional to cost alone are shown at the right of the figures. These figures also illustrate the differences in decision making speeds among students since both figures represent the total six hours of training.

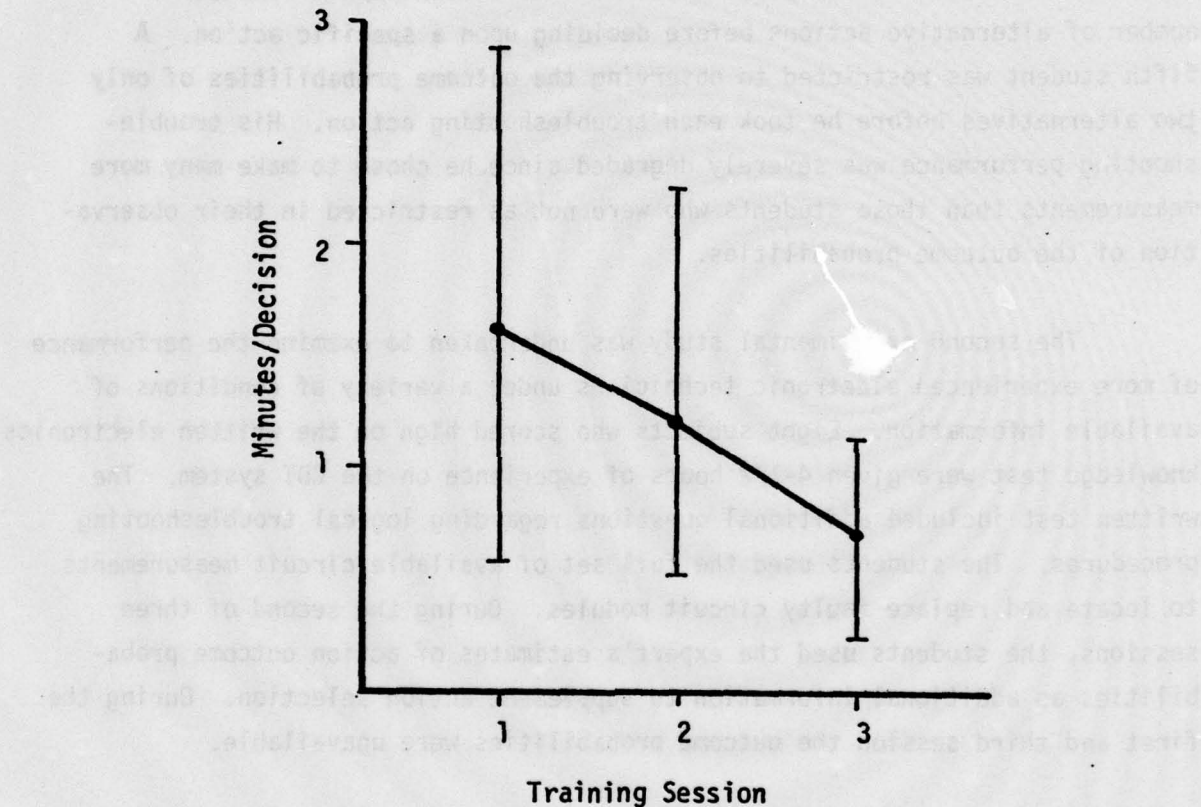
The experiment also demonstrated that the students rely heavily on the presentation of action probabilities. The student would consider a number of alternative actions before deciding upon a specific action. A fifth student was restricted to observing the outcome probabilities of only two alternatives before he took each troubleshooting action. His troubleshooting performance was severely degraded since he chose to make many more measurements than those students who were not as restricted in their observation of the outcome probabilities.

The second experimental study was undertaken to examine the performance of more experienced electronic technicians under a variety of conditions of available information. Eight subjects who scored high on the written electronics knowledge test were given 4-1/2 hours of experience on the CDT system. The written test included additional questions regarding logical troubleshooting procedures. The students used the full set of available circuit measurements to locate and replace faulty circuit modules. During the second of three sessions, the students used the expert's estimates of action outcome probabilities as additional information to supplement action selection. During the first and third session the outcome probabilities were unavailable.

The data indicate that the students continued to improve their decision making speed throughout the three training sessions. The mean and range of decision time performance for five of the eight scheduled subjects are shown in Figure 4-7. As shown in Figure 4-8, when the outcome probabilities are withdrawn in session 3 the students' decision efficiency decreases.

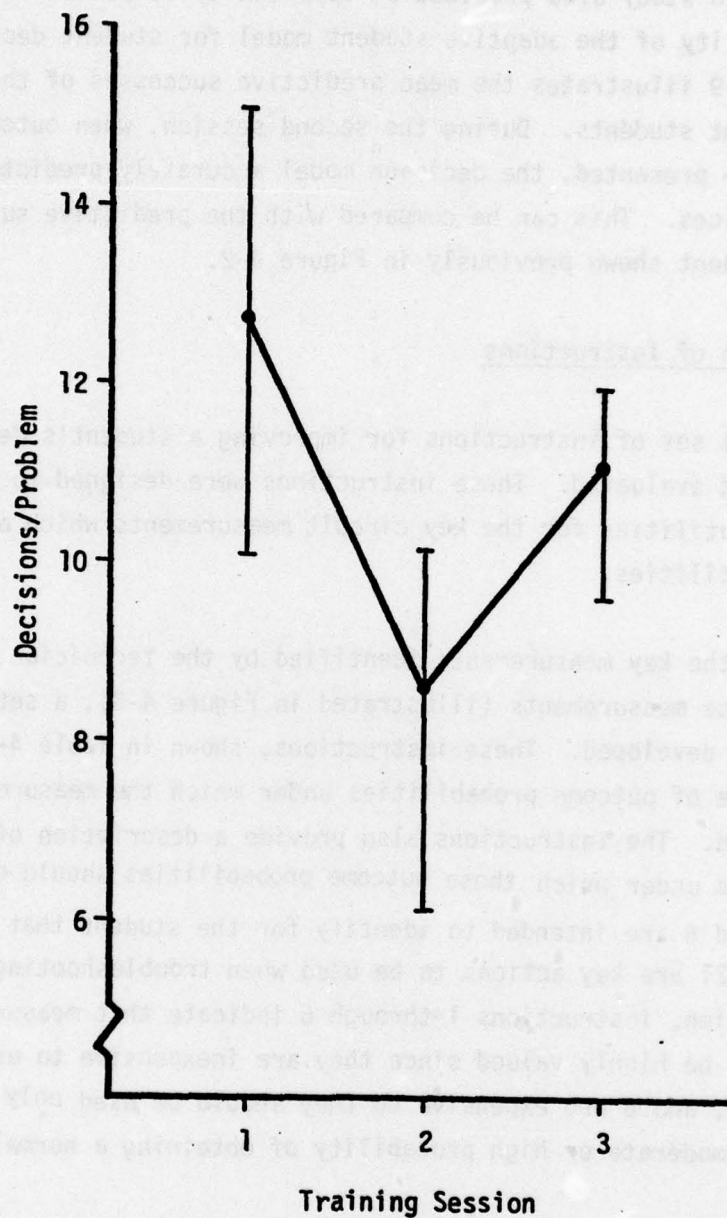
calculated to be proportional to cost alone are shown at the right of the figures. These figures also illustrate the differences in decision making speeds among students since both figures represent the total six hours of training.

The experiment also demonstrated that the students rely heavily on the presentation of actual probabilities. The student would consider a



**STUDENT DECISION TIME AS
A FUNCTION OF TRAINING**

Figure 4-7



STUDENT DECISION EFFICIENCY AS
A FUNCTION OF TRAINING

Figure 4-8

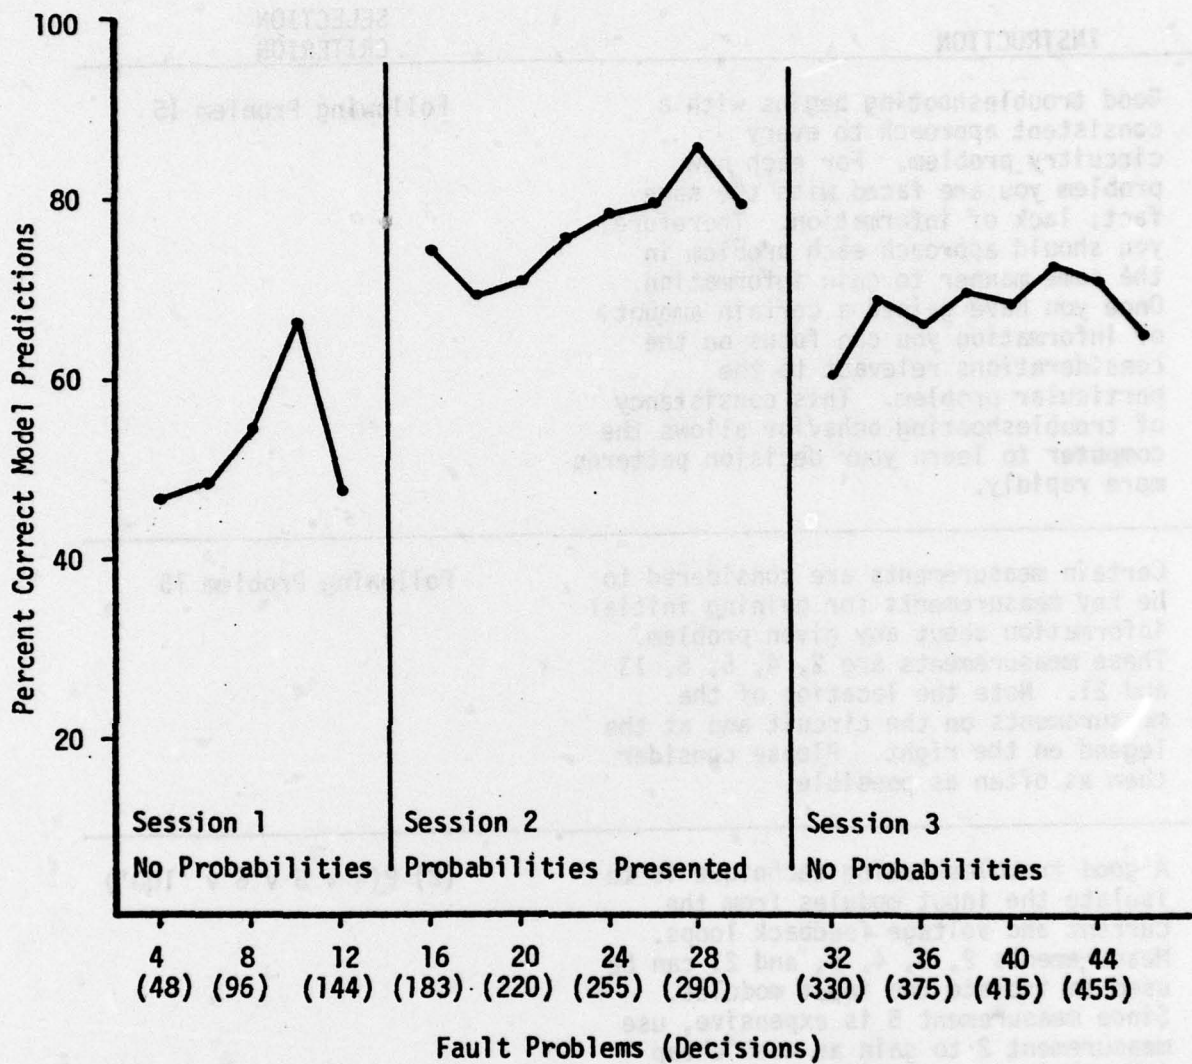
This second study also provided an opportunity to assess the predictive capability of the adaptive student model for student decision makers. Figure 4-9 illustrates the mean predictive successes of the student model for the eight students. During the second session, when outcome probabilities were presented, the decision model accurately predicted 75% of the student's choices. This can be compared with the predictive success with the simulated student shown previously in Figure 4-2.

4.3 Evaluation of Instructions

An initial set of instructions for improving a student's decisions were developed and evaluated. These instructions were designed to teach the student to adopt utilities for the key circuit measurements which are similar to the expert's utilities.

Based on the key measurements identified by the technician and his utilities for these measurements (illustrated in Figure 4-3), a set of six instructions were developed. These instructions, shown in Table 4-1, identify the range of outcome probabilities under which the measurements should be selected. The instructions also provide a description of the circuit conditions under which those outcome probabilities should occur. Instructions A and B are intended to identify for the student that measurements 2, 4, 5, 6, 13, and 21 are key actions to be used when troubleshooting the power supply. In addition, instructions 1 through 6 indicate that measurements 2, 13, and 21 should be highly valued since they are inexpensive to use while measurements 4, 5, and 6 are expensive so they should be used only in situations of moderate or high probability of obtaining a normal measurement outcome.

Table 4-1 also shows the criteria for selecting the instructions. These criteria are shown as Boolean expressions, using the operators and relations described in Section 3.6. An additional operator was added to



PREDICTIVE SUCCESS OF DECISION MODEL

DURING STUDENT TRAINING

Figure 4-9

TABLE 4-1. UTILITY TRAINING INSTRUCTIONS
AND SELECTION CRITERIA

INSTRUCTION	SELECTION CRITERION
A. Good troubleshooting begins with a consistent approach to every circuitry problem. For each new problem you are faced with the same fact; lack of information. Therefore, you should approach each problem in the same manner to gain information. Once you have gained a certain amount of information you can focus on the considerations relevant to the particular problem. This consistency of troubleshooting behavior allows the computer to learn your decision patterns more rapidly.	Following Problem 15
B. Certain measurements are considered to be key measurements for gaining initial information about any given problem. These measurements are 2, 4, 5, 6, 13 and 21. Note the location of the measurements on the circuit and at the legend on the right. Please consider them as often as possible.	Following Problem 15
1. A good troubleshooting technique is to isolate the input modules from the current and voltage feedback loops. Measurements 2, 3, 4, 5, and 21 can be used to isolate the input modules. Since measurement 5 is expensive, use measurement 2 to gain as much cheap information as possible. Use measurement 2 when it has a low probability of a normal outcome to confirm the state of the transformer.	(2) P(4 v 5 v 6 v '100')

TABLE 4-1. (CONTINUED)

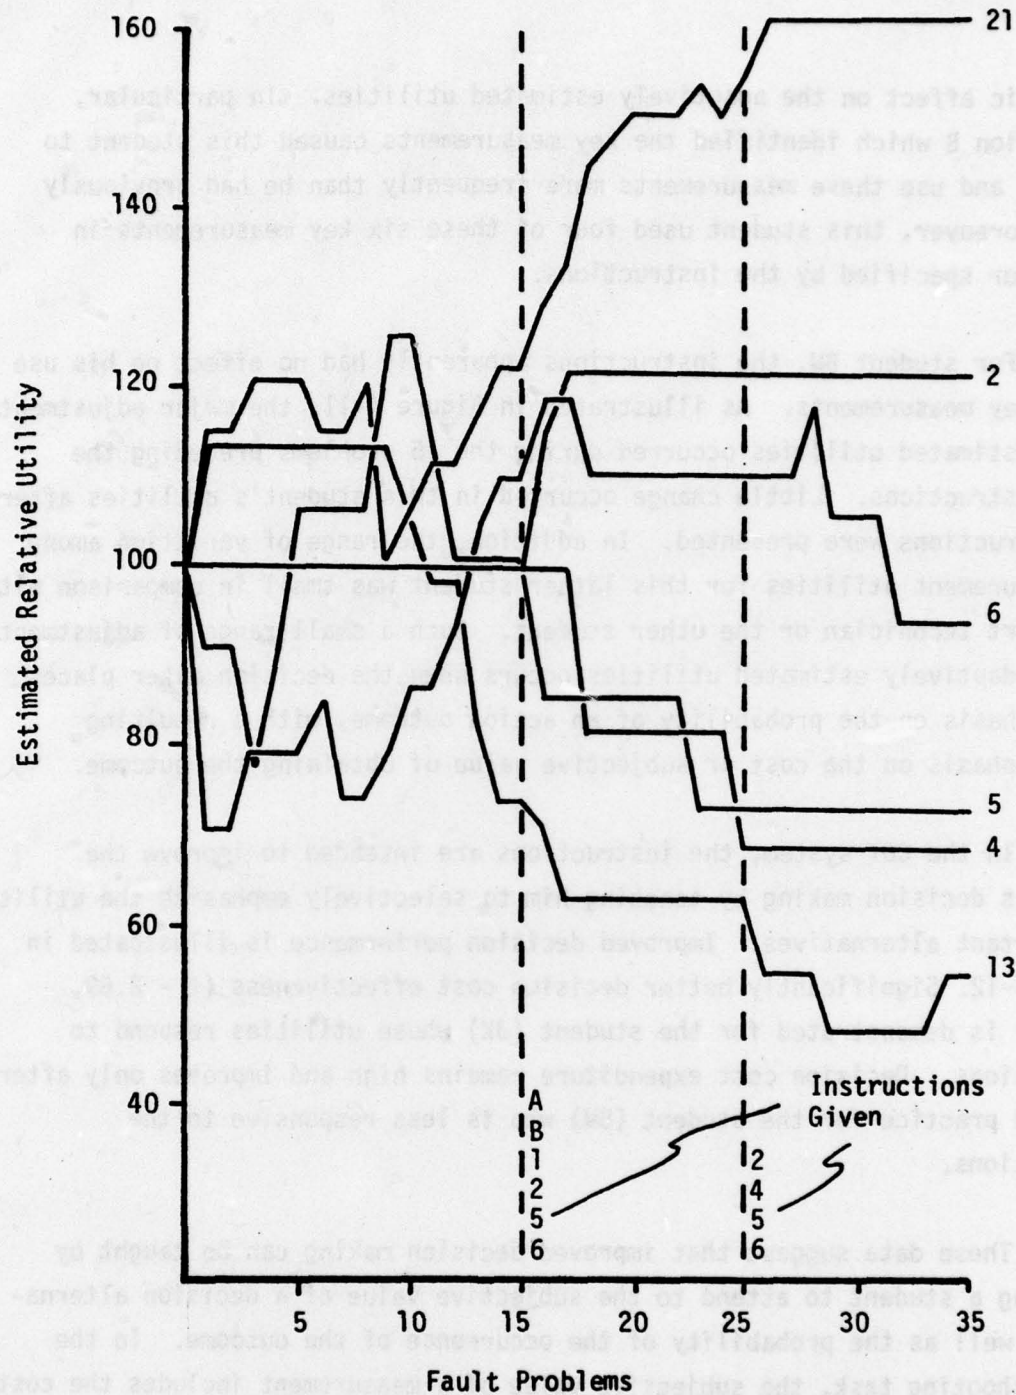
INSTRUCTION	SELECTION CRITERION
2. A good first step in checking the operation of the current and voltage feedback loops is to check the output of the series regulator. This should be done with the circuit operating at full output since this fully exercises the circuit functions. Therefore, measurement 11 or 13 should be used even if there is a low probability of a normal outcome. Use measurement 13 since it is much cheaper than 11.	(13) P (4 v 5 v 6 v '100')
3. Measurement 21 serves well to isolate the operation of the transformer and reference DC source. Since this measurement is relatively inexpensive, use it even when the probability of the normal outcome is low.	(21) P (4 v 5 v 6 v '100')
4. Measurement 4 is required to isolate the transformer from the reference DC source. However, measurement 4 is relatively expensive so it should be used only after a non-normal outcome from measurement 21 has been found. In this case the table will show a probability of about 50% for the normal outcome of measurement 4.	(2 v 13 v 21) P (4) Δ ('100') P (4)
5. Although Measurement 5 is located at a good point to isolate the power input modules, it is expensive. Use this measurement after you have eliminated most other possibilities. Measurement 5 should be used when the probability of a normal outcome is rather high but not certain (a range of 60% to 80%).	(2 v 13 v 21 v '80') P (5)
6. Measurement 6 should be used to determine the status of the current feedback loops. However, it is expensive so use it only after finding that Measurements 13, 19, or 20 are not normal. In this case, the probability of a normal outcome for measurement 6 will be about 60% to 80%.	(2 v 13 v 21 v '80') P (6)

enote an absolute utility value, e.g., '100', in addition to an adaptively-estimated utility for a specific measurement.

Two students who scored high on the written test of electronics knowledge participated in this initial evaluation of the instructions. During the first session, which lasted about two hours, each student read the written introduction to the system and then practiced with five troubleshooting problems on the CRT display terminal. The written introduction (1) described how to operate the CDT system, (2) discussed the operation of the power supply circuit and its functional modules, and (3) gave a brief introduction to the notion of probabilistic outcomes to alternatives. However, no instructions about possible decision strategies or alternatives in troubleshooting the power supply were given in this introduction.

During the second training session, each student completed 15 troubleshooting problems at the display terminal. No instructions were given during this session; however, the adaptive model of the student generated the estimates of the student's utilities for measurement outcomes. At the beginning of the third training session (preceding problem 16), selected instructions from the instructions shown in Table 4.1 were given to the student. Each instruction was typed on an index card to which he could refer at any time during the training session. The two general instructions (A and B) were given plus any instructions for which the associated criterion was TRUE. This third session continued until the student had completed problems 16 through 25. A fourth training session included problems 26 through 35. At the beginning of this final session, all instructions were given for which the associated criteria were TRUE.

The results of this evaluation indicate that the effects of the instructions are not the same for different students. As illustrated in Figure 4-10, the instructions given to student JK following problem 15 had



RELATIVE UTILITY FOR KEY MEASUREMENTS
AS A FUNCTION OF FAULT PROBLEMS
(STUDENT JK)

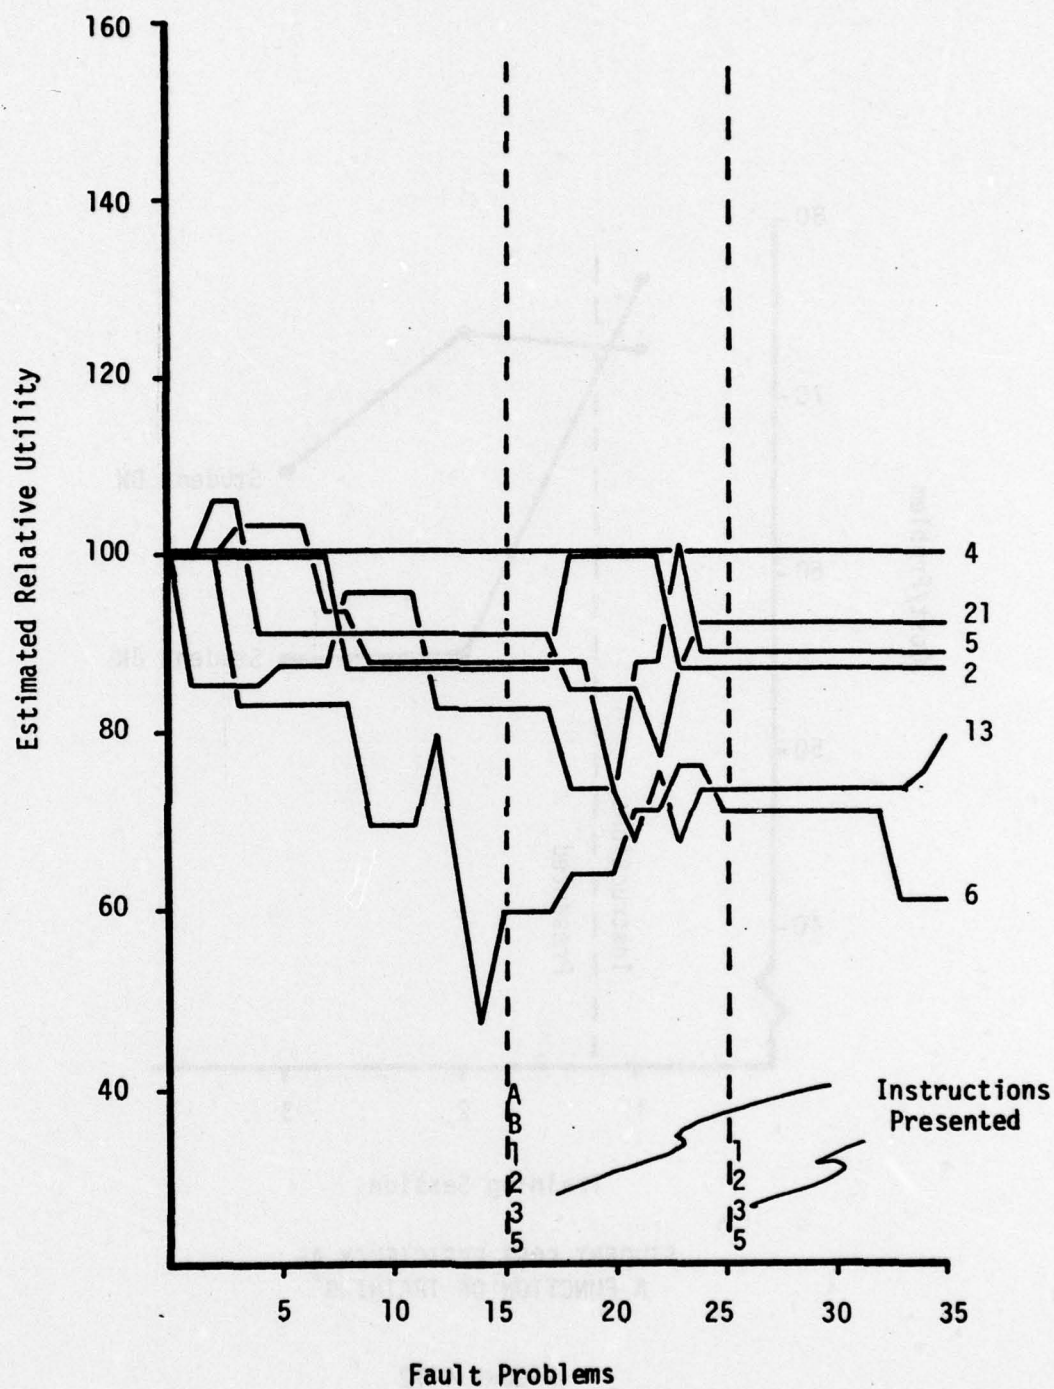
FIGURE 4-10

a dramatic effect on the adaptively estimated utilities. In particular, instruction B which identified the key measurements caused this student to consider and use these measurements more frequently than he had previously done. Moreover, this student used four of these six key measurements in the manner specified by the instructions.

For student BW, the instructions apparently had no effect on his use of the key measurements. As illustrated in Figure 4-11, the major adjustments in the estimated utilities occurred during the 15 problems preceding the first instructions. Little change occurred in this student's utilities after the instructions were presented. In addition, the range of variation among all measurement utilities for this latter student was small in comparison with the expert technician or the other student. Such a small range of adjustment in the adaptively estimated utilities occurs when the decision maker places high emphasis on the probability of an action outcome, with a resulting lower emphasis on the cost or subjective value of obtaining the outcome.

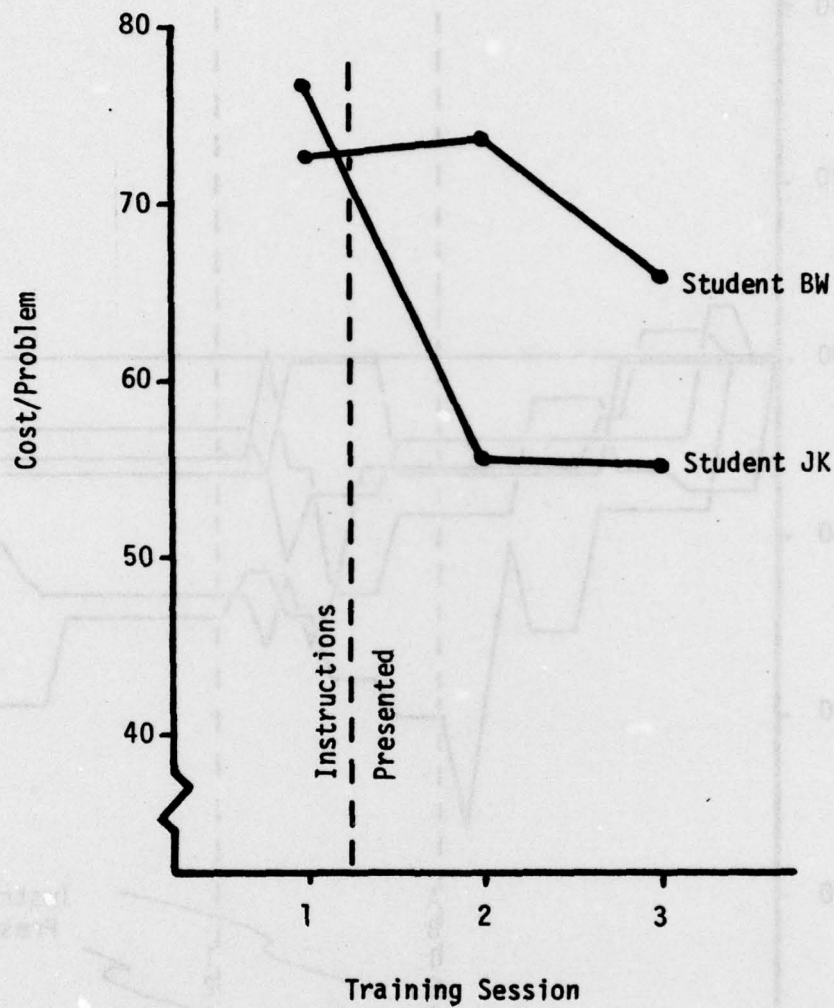
In the CDT system, the instructions are intended to improve the student's decision making by teaching him to selectively emphasize the utility of important alternatives. Improved decision performance is illustrated in Figure 4-12. Significantly better decision cost effectiveness ($t = 2.69$, $p \leq .01$) is demonstrated for the student (JK) whose utilities respond to instructions. Decision cost expenditure remains high and improves only after extended practice for the student (BW) who is less responsive to the instructions.

These data suggest that improved decision making can be taught by directing a student to attend to the subjective value of a decision alternative as well as the probability of the occurrence of the outcome. In the troubleshooting task, the subjective value of a measurement includes the cost of making the measurement (expressed in terms of time and money), the effectiveness of the measurement in isolating a malfunction, the functional or topological location of the measurement point in the circuit, etc.



RELATIVE UTILITY FOR KEY MEASUREMENTS
AS A FUNCTION OF FAULT PROBLEMS
(STUDENT BW)

FIGURE 4-11



STUDENT COST EFFICIENCY AS
A FUNCTION OF TRAINING

FIGURE 4-12

However, the utility of a decision alternative must be sufficiently salient to the decision maker to have a significant effect on his decisions. In the present CDT system, the student's utilities are a function of the stated costs and his knowledge of the effectiveness of the various measurements. The performance data from student BW suggest that additional reinforcement is required to emphasize the utility of each decision alternative.

4.4 Discussion

The current year's efforts have demonstrated that the adaptive expected utility model tracks the performance of consistent decision makers. In addition, comparison of an expert's stated ranking of measurement importance with his adaptively-estimated utilities has shown that the adaptive model ranks his preferences accurately. Thus, the ability of the model to reproduce human behavior within the context of the troubleshooting task has been demonstrated. Experimental sessions with experienced and inexperienced student technicians have also demonstrated that the simulated circuit model provides an accurate representation of circuit troubleshooting and that students demonstrate improved decision making performance with extended practice.

The demonstration of the validity of the adaptive model and of the accuracy of the circuit simulation provides a firm background for developing the instructional feedback to train the students in evaluating their decision alternatives. An initial set of instructions were developed which focus a student's attention on the proper utilization of the key circuit measurements. A limited evaluation has demonstrated that these instructions are effective with a student who has demonstrated a sensitivity to utilities for alternatives. This limited study suggests that explicit reinforcement and immediate knowledge of results will be required to effectively alter all students' performance. Further development of the training capabilities of the CDT system should incorporate such explicit reinforcement and additional feedback algorithms designed to strengthen particular deficiencies identified in students' decision making.

5. REFERENCES

1. Baker, J. D. Quantitative Modeling of Human Performance in Information Systems. *Ergonomics*, 1970, 13:645-644.
2. Baker, J. D. SIMTOS: A Man-in-the-Loop Interactive Simulation of a Tactical Operations System. Paper presented at the NATO Program Panel on Systems Science, Munich, Germany, August 1974.
3. Beach, B. H. Direct and Indirect Methods for Measuring Utility. Department of Psychology, University of Washington, Scientific Interim Report, July 4, 1972.
4. Bond, N. A. and Rigney, J. W. Bayesian Aspects of Troubleshooting Behavior, *Human Factors*, 1966, 8(5):377-383.
5. Brown, R. V., Hoblitzell, C. M., Peterson, C. R., and Ulvila, J. W. Decision and Designs, Inc., Final Technical Report No. 1, No. 74-2, N00014-74-C-0263, Code 455, Office of Naval Research, September 1974.
6. Coombs, C. H. and Pruitt, D. G. Components of Risk in Decision Making: Probability and Variance Preferences. *Journal of Experimental Psychology*, 1960, 60:265-277.
7. Davis, K. B., Weisbrod, R. L., Freedy, A., and Weltman, G. Adaptive Computer Aiding in Dynamic Decision Processes: An Experimental Study of Aiding Effectiveness. Perceptronics, Inc., Woodland Hills, CA, Technical Report PTR-1016-75-5, May 1975.
8. Dawes, R. M. Graduate Admission: A Case Study. Oregon Research Institute, Technical Report 10(1), 1970.
9. Edwards, W. Dynamic Decision Theory and Probabilistic Information Processing, *Human Factors*, 1962, 4:59-74.
10. Fishburn, P. C. *Decision and Value Theory*, New York: Wiley, 1964.
11. Fishburn, P. C. Methods of Estimating Additive Utilities. *Management Science*, March 1967, 13(7):435-453.

12. Freedy, A., Weisbrod, R., Davis, K., May, D., and Weltman, G. Adaptive Computer Aiding in Dynamic Decision Processes: Adaptive Decision Models and Dynamic Utility Estimation. Perceptronics, Woodland Hills, CA, Technical Report PTR-1016-74-5(I), May 1974.
13. Goodman, B., Saltzman, M., Edwards, W., and Drantz, D. M. When SEU is Irrelevant. University of Michigan, Dept. of Psychology, 1971.
14. Halpin, S. M., Thronberry, J. A., and Streufert, S. The Credibility of Computer Estimates in a Simple Decision Making Task. Purdue University, Technical Report No. 5, Contract N00014-67-A-0226-0020, January 1973.
15. Hanes, R. M. and Gebhard, J. W. The Computer's Role in Command Decision. U.S. Naval Institute Proceedings, September 1966.
16. Howard, Ronald A. The Foundations of Decision Analysis. IEEE Transactions on Systems Science and Cybernetics, September 1968, SSC-4(3):211-219.
17. Kelly, C. W. and Peterson, C. R. Probability Estimates and Probabilistic Procedures in Current Intelligence Analysis. Federal Systems Division, IBM Corp., Report FSG 5047, January 1971.
18. Kneppreth, N. P., Gustafson, D. H., Johnson, E. M., and Leifer, P. Techniques for the Assessment of Worth. U.S. Army Research Institute for the Behavioral and Social Sciences, Technical Paper 254, Contract DAHC19-72-C-0026, December 1973.
19. Krantz, D. C., Luce, R. D., Suppes, P., and Tversky, A. Foundations of Measurement: Additive and Polynomical Representations, I. New York: Academic Press, 1971.
20. Lichtenstein, S. and Slovic, P. Reversals of Preferences Between Bids and Choices in Gambling Decisions. Journal of Experimental Psychology, 1971, 89:46-55.
21. Luce, R. D. and Suppes, P. Preference, Utility, and Subjective Probability. In Handbook of Mathematical Psychology, 42-49, Luce, R. D., Bush, R. R., and Galanter, E. (Eds.) New York: Wiley, 1965.
22. Luce, R. D. and Raiffa, H. Games and Decisions: Introduction and Critical Survey, New York: Wiley, 1957.

23. May, D. M., Crooks, W. H., Purcell, D. D., Lucaccini, L. F., Freedy, A., and Weltman, G. Application of Adaptive Decision Aiding Systems to Computer-Assisted Instruction. Perceptronics, Inc. Woodland Hills, CA, Technical Report PATR-1019-74-12, December 1974.
24. Miller, L. W., Kaplan, R. J., and Edwards, W. JUDGE: A Value-Judgment-Based Tactical Command System. Organizational Behavior and Human Performance, 1967, 2:329-374.
25. Nickerson, R. S. and Feehrer, C. E. Decision Making and Training: A Review of Theoretical and Empirical Studies of Decision Making and Their Implications for the Training of Decision Makers. Bolt, Beranek and Newman, Cambridge, MA, Technical Report NAVTRAEQUIPCEN 73-C-0128-1, July 1975.
26. Nilsson, N. J. Learning Machines. New York: McGraw, 1965.
27. Payne, J. R., Miller, A. C., and Roney, J. V. The Naval Task Force Decision Environment. Stanford Research Institute, Technical Report NWRC-TR-8, SRI Project No. 3315, N00014-74-C-0357/NR 274-199, Office of Naval Research, September 1974.
28. Shannon, C. E. and Weaver, W. The Mathematical Theory of Communication. Urbana: University of Illinois Press, 1949.
29. Steeb, R., Artof, M., Crooks, W. H., Freedy, A., and Weltman, G. Man/Machine Interaction in Adaptive Computer-Aided Control: Final Report. Perceptronics, Woodland Hills, CA, Technical Report PATR-1008-75-12/31, December 1975.
30. Tversky, A. Additivity, Utility, and Subjective-Probability. Journal of Mathematical Psychology, 1967, 4:175-202.
31. Vaughn, W. S. (Jr.) and Mavor, A. S. Behavioral Characteristics of Men in the Performance of Some Decision-Making Task Components. Ergonomics, 1972.
32. Weisbrod, R. L., Davis, K. B., Freedy, A., and Weltman, G. Adaptive Computer Aiding in Dynamic Decision Processes. Perceptronics, Woodland Hills, CA, Technical Report PTR-1016-74-11, November 29, 1974.
33. Weisbrod, R. L., Davis, K. B., and Freedy, A. Adaptive Utility Assessment in Dynamic Decision Processes: An Experimental Evaluation of Decision Aiding. Proc. International Conf. on Systems, Man, and Cybernetics, San Francisco, CA, September 1975.
34. Wendt, D. Some Criticisms of the General Models Used in Decision Making Experiments. Engineering Psychology Laboratory (University of Michigan) Technical Report 011313-8-T, November 1973.
35. Wendt, D. Utility and Risk. Acta Psychologica, 1970, 34:214-228.